Computer Science and Artificial Intelligence Laboratory
Technical Report

MIT-CSAIL-TR-2011-024 April 21, 2011

Mechanism Design with Approximate Valuations

Alessandro Chiesa, Silvio Micali, and Zeyuan Allen Zhu

massachusetts institute of technology, cambridge, ma 02139 usa — www.csail.mit.edu

Mechanism Design with Approximate Valuations

Alessandro Chiesa[*] MIT CSAIL
Silvio Micali[†] MIT CSAIL
Zeyuan Allen Zhu[‡] MIT CSAIL

April 8, 2011

Abstract

In mechanism design, we replace the strong assumption that each player knows his own payoff type *exactly* with the more realistic assumption that he knows it only *approximately*. Specifically, we study the classical problem of maximizing social welfare in single-good auctions when players know their true valuations only within a constant multiplicative factor $\delta \in (0, 1)$.

Our approach is deliberately non-Bayesian and very conservative: each player i only knows that his true valuation is one among finitely many values in a δ*-approximate set* K_i, and his true valuation is *adversarially and secretly* chosen in K_i at the beginning of the auction.

We prove tight upper and lower bounds for the fraction of the maximum social welfare achievable in our model, in either dominant or undominated strategies, both via deterministic and probabilistic mechanisms. The landscape emerging is quite unusual and intriguing.

Prior Version: Technical Report MIT-CSAIL-TR-2011-009, MIT, February 2011.

[*]Email: alexch@csail.mit.edu.
[†]Email: silvio@csail.mit.edu.
[‡]Email: zeyuan@csail.mit.edu.

1 Introduction

Mechanism design aims at leveraging the players' knowledge and rationality, so as to produce desirable outcomes, based on the assumption that each player knows his own true pay-off type. We raise a natural and general question:

What happens to mechanism design when players **approximately** *know their own pay-off types?*

1.1 The Approximate-Valuation Model

We could define our approximate-valuation model, together with its basic notions (of dominance, implementations, etc.), in a very abstract setting. For simplicity, however, in this paper we define it just for auctions of a single good.

A realistic motivation. The possibility that a player may have *approximate*, rather than *exact*, knowledge of his own valuation strikes us to be quite realistic. For instance, consider a firm about to participate in an auction of a given good. In such an auction, a player's *valuation* consists of a non-negative number representing the player's value for the good for sale. For finiteness sake, we assume that all valuations are integers in the interval $[0, B]$, where B is a suitably large (upper)bound. Then, no one would be too surprised if different employees, when asked to figure out the firm's *true* valuation for the good, reported different answers; nor if a given employee reported a range of values rather than a single value.

A set-theoretic approach to approximation. Our model for the knowledge that each player i has about his own true valuation tv_i is *set theoretic* (and is inspired by [CM10]). Specifically, we assume that there is some constant $\delta \in [0, 1)$ such that each player i knows a "δ-approximate" subset of $\{0, 1, \ldots, B\}$ containing tv_i. That is, even though player i does not know tv_i, player i knows a subset K_i of $\{0, \ldots, B\}$, which we call the *approximate-valuation set* of player i, such that:

$$\text{(i) } \mathsf{tv}_i \in K_i \quad \text{ and } \quad \text{(ii) } K_i \subseteq \big[(1-\delta)x_i, (1+\delta)x_i\big] \cap \mathbb{Z} \text{ for some value } x_i \in \mathbb{R}.$$

Essentially, the *approximation accuracy* δ "upper bounds the coarseness of each player's knowledge about his own valuation". For instance, if $\delta = 0.1$, then everyone knows that each player knows his own true valuation within a 10% accuracy. We remark that:

- To achieve a greater level of generality, we may consider two distinct accuracy parameters: a multiplicative one, δ^*, and an additive one, δ^+, leading to the following modified constraint:

  $$K_i \subseteq \big[(1-\delta^*)x_i - \delta^+, (1+\delta^*)x_i + \delta^+\big] \cap \mathbb{Z} \text{ for some value } x_i \in \mathbb{R}.$$

  For simplicity, however, in this extended abstract we consider a single accuracy parameter, and we find the multiplicative one more meaningful.

- The approximate-valuation set K_i can be an arbitrary subset of the integers in the interval $[(1-\delta)x_i, (1+\delta)x_i]$, and in principle could consist of all the primes in it. For example, K_i may consist of just two values. Consider a player i who is about to participate to a yard-sale auction of a large piece of furniture that would be illegal for him to transport on top of his car. Then, even if he knew exactly his valuation for the piece, he should consider "subtracting" the fine he might incur on the way home. Some players will approach this valuation uncertainty by computing the probability of being caught by the police, based on the time of the day, the traffic pattern, the probability that the police has bigger crimes to worry about that day, and so on. But some others may prefer taking a set-theoretic approach and forgo the above complexities —after all computation too has a cost!

- The approximation accuracy is a *global* parameter, known to the designer. The strategic behavior of an individual player i is determined by his own set K_i, which is not known to the designer. While δ-approximate, K_i may consist of fewer and well-clustered possible valuations, so that in the end, player i de facto has a smaller "individual" approximation accuracy.
- Valuations may actually be very approximate indeed. Consider a firm participating to an auction for the exclusive rights to manufacture solar panels in the US for a period of ten years. Even if the demand were precisely known in advance, and the only uncertainty were to come from the firm's ability to lower its costs of production via some breakthrough research, an approximation accuracy of the firm's own valuation for the license could easily exceed 0.5.

Interpretation. Ours is a very adversarial model. The exact true valuation tv_i of player i, while promised to lie in K_i should be thought as having being selected *secretly* (even with respect to i himself) by an *adversary* at the start of the auction. Before bidding, a player i only knows his knowledge set K_i, and thus that, in an outcome in which he pays P_i, his utility is $-P_i$ if he does not win the good, and a value in the set $\{v - P_i : v \in K_i\}$ otherwise. But, in the latter case, he may never realize his *exact* utility, $\mathsf{tv}_i - P_i$, or he may realize it only after bidding, or after the auction is over —e.g., a day or a year later. (This may sound strange, but is actually quite reasonable, in auctions and other strategic situations. For instance, many colonial powers realized to have received negative utility only after their colonies —fortunately— regained their freedom.)

Our definitions of dominance will thus be defined also in an adversarial way. Informally speaking, a strategy s will be more favored than a strategy t, only when for all possible valuations inside his knowledge set K_i, s is better than t. By investigating a model of an adversarial nature we may be able to prove less, but our results will then be guaranteed to hold for a larger variety of more "realistic" settings.

Difference from Bayesian. We emphasize that our approximate-valuation model is "safer" than assuming that i knows an "individual Bayesian" D_i, from which his true valuation has been drawn. Indeed, "i knows D_i" is a very strong assumption in an auction of a single good: player i may behave as if his true valuation were *exactly* the expectation of D_i, reducing the problem into the exact-valuation world. Instead, if i knows that his valuation is in the support of some (unknown) distribution D_i, setting this support to be the set K_i, he does not have such a privilege.

1.2 A Dual View of Our Model

Our *set-theoretic* model can however be re-interpreted as an *approximate-Bayesian* model. In particular, whether or not the entire payoff-type profile is drawn from a given prior (and whether or not some information about this prior may be known to the designer or common knowledge to the players), a player i may *individually* know that his own true valuation is drawn from some distribution D_i in some class of distributions $\mathscr{D}_i$, without being sure of *which* distribution in $\mathscr{D}_i$ is the right one. In such a setting, "$\mathscr{D}_i$ plays the role of K_i".

Beyond such a generic parallelism, it is actually possible to convert the above Bayesian model into ours, so that our results and mechanisms apply. For example, suppose that each player i knows that, for each possible valuation v, the probability that his true valuation is v is between $(1-\delta)p_i(v)$ and $(1+\delta)p_i(v)$; then, each $\mathscr{D}_i$ is the class of all distributions that are consistent with player i's probability constraints, and it is easy to show that the expectations of distributions in $\mathscr{D}_i$ form a δ-approximate set K_i in our set-theoretic model.

A more thorough investigation of this "duality" will however be conducted separately.

1.3 The First Goal

For concreteness, we choose to investigate the feasibility of mechanism design in the approximate-valuation model for a very simple and familiar application, *maximizing social welfare in single-good auctions,* for which the second-price mechanism gives us a simple, elegant, and perfect solution in the traditional, exact-valuation model. That is, we wish to start the study of approximate-valuation mechanism design by providing an answer to the following question:

How much social welfare can we guarantee in approximate-valuation auctions?

1.4 Results

Since we are in a non-Bayesian setting, but in the presence of substantial uncertainty, two types of implementations are natural to explore: in dominant strategies and in undominated strategies.

1.4.1 The Inadequacy of Dominant-Strategy Mechanisms

When the players know exactly their true valuations, the second-price mechanism is universally known to maximize social welfare in dominant strategies. A striking signal that the approximate-valuation model is a "new world" is that dominant strategies, though a priori meaningful, now become useless.

Superficially, this may seem obvious: if you are uncertain whether your valuation is v or v', how can you have a "best bid"? Less superficially, we have to consider that, when the players have approximate-valuation sets rather than exact valuations, then a reasonable mechanism should give them the option to report sets of valuations (rather than single valuations). Thus, among such mechanisms there might be a dominant-strategy one that maximizes social welfare. Perhaps more realistically, one may expect some degradation of performance due to the approximate accuracy of the players' self knowledge. A priori, it would be legitimate to conjecture that under the assumption that each player knows his own valuation with accuracy $\delta \in (0, 1)$, a dominant strategy mechanism might guarantee only a fraction of the underlying, true maximum social welfare: for instance, a fraction $(1 - \delta)$, $(1 - 3\delta)$, or $(1 - \delta)^2$. We prove, however, that this is *not at all* the case. Namely,

Theorem 1 (informal). *The best fraction of the maximum social welfare that can be guaranteed by a (possibly probabilistic) dominant-strategy mechanism in the approximate-valuation model is*

$$\approx \frac{1}{n} \ .$$

Note that $\frac{1}{n}$ of the social welfare can be trivially achieved by the dominant-strategy mechanism that ignores all bids and assigns the good to a random player! Thus our theorem says that, in the approximate-valuation model, no matter how small the approximation accuracy may be, there is nothing smarter to do in dominant strategies. Thus, one way to interpret our result is the following:

Dominant strategies are **intrinsically** *linked to the* **exact** *knowledge of our own valuations.*

When at least one player's valuation is known to have positive approximation accuracy (e.g., he knows his own valuation only within 0.001%), then, as long as the possible valuations are sufficiently (in the previous case, that would be 2,000), *demanding the use of a dominant-strategy mechanism to generate high social welfare is tantamount to demanding that the good be assigned at random.*

New worlds, new realities.

1.4.2 The Power of Deterministic Undominated-Strategy Mechanisms

In light of Theorem 1, implementation in undominated strategies now becomes a natural choice. For such type of implementation, we prove the following theorem. It states that even deterministic, undominated-strategy mechanisms can guarantee a good fraction of the social welfare, that is, a fraction that depends solely on the approximation accuracy δ. Indeed, the second-price mechanism does this job, and is optimally too. Namely,

Theorem 2 (**informal**). *In the approximate-valuation model, the best fraction of the maximum social welfare that can be guaranteed in undominated strategies by a **deterministic** mechanism is*

$$\approx \left(\frac{1-\delta}{1+\delta}\right)^2 .$$

Moreover, the second-price mechanism with lexicographic tie-breaking rule essentially guarantees such performance.

That is, in the approximate-valuation world, although the second-price mechanism is useless under the solution concept of dominant strategies (since it does not provide any dominant strategy), it gives modest social welfare guarantee in undominated strategies. (The second-price mechanism actually delivers excellent social welfare in undominated strategies also in the exact-valuation world, but this fact is overshadowed by its perfect performance in dominant strategies.)

In the exact-valuation world, proving that the second-price mechanism maximizes social welfare in dominant strategies is almost immediate. However, proving Theorem 2 is more involved. To begin with, any analysis in undominated strategies is more complex. In addition, in the approximate-valuation world this difficulty is compounded by the fact that each player has to consider many candidates for his own valuation rather than a single one. We thus choose to prove our theorem only after establishing some structural properties of the new world.

1.4.3 The Greater Power of Probabilistic Undominated-Strategy Mechanisms

We prove that, there exists a performance gap between probabilistic and deterministic mechanisms working in undominated strategies. Namely,

Theorem 3 (**informal**). *For every $\delta \in (0, 1)$, there exists a **probabilistic**, undominated-strategy mechanism M_{opt} satisfying the following two properties in the approximate-valuation model:*

- *Its guaranteed fraction of the maximum social welfare is*

$$\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} .$$

- *Its guarantee is essentially optimal among all undominated-strategy mechanisms.*

The performance gap between probabilistic and deterministic implementations in undominated strategies tends to zero as the number of players increases, but is quite relevant in concrete scenarios. For instance, when $\delta = 0.5$, our mechanism M_{opt} guarantees a social welfare that is at least **five** times higher that of the second-price mechanism when there are 2 players, and at least **three** times higher when there are 4 players. Even when $\delta = 0.25$, the guaranteed performance of our mechanism is almost **two** times higher than that of the second-price when there are 2 players. (See Appendix F for a chart comparing our optimal mechanism with the second-price one.)

Besides its potentially practical relevance, Theorem 3 is actually of theoretical value. Indeed, it emphasizes another difference between the exact-valuation and approximate-valuation worlds. In the exact-valuation world, no performance gap like ours exists between the probabilistic and deterministic mechanisms working in undominated strategies. (Due to the excellent performance of the deterministic second-price mechanism, the only possible gap is a mere additive 1.) In addition, the proof of Theorem 3 has required us to develop a set of techniques that we believe to be of independent interest.

1.5 Techniques

Exploring a new direction requires developing a new set of tools. To begin with, extending to the approximate-valuation world even the basic concept of weak dominance and the revelation principle requires some care. (Else, one runs the risk of "defining out" crucial alternatives and forcing impossibilities.) In addition, to guide the design and analysis of mechanisms in our new setting, we find it useful to establish two structural lemmas. Let us describe the simpler one first.

The Undominated Intersection Lemma. To prove impossibility results for general mechanisms, we must deal with arbitrary sets of strategies. In our approximate-valuation model, even the "natural" strategies are richer than before. For instance, a player i having an approximate-valuation set K_i might report a single number (e.g., a member in K_i), or a set of numbers (e.g., K_i itself). However, no matter what strategies a mechanism may grant, we prove that for any two sufficiently overlapping approximate-valuation sets K_i and K_i', the corresponding sets of undominated strategies S and S' must have a pair of (mixed) strategies that are "arbitrarily close".

This lemma has been key to establish our impossibility results for undominated strategies.

The Distinguishable Monotonicity Lemma. To prove possibility results, we are instead happy to consider only special sets of strategies. In particular, we consider mechanisms that constrain the players to report only individual numbers in $\{0, 1, \ldots, B\}$, like direct mechanisms in the exact-valuation model. In both cases, a player's strategies are well-ordered, and it is thus meaningful to ask whether his probability of getting the good grows monotonically with his reports (keeping the ones of his opponents fixed). Indeed, a proper monotonicity condition characterizes which single-good auction mechanisms are dominant-strategy in the exact-valuation model. The advantage of artificially restricting the strategies to $\{0, 1, \ldots, B\}$ in the approximate-valuation model is that a simple lemma, based on a slightly stronger monotonicity condition, provides a clean characterization of the set of *undominated strategies* associated to any approximate-valuation set K_i of a player i: namely, the strategies between the minimum integer and the maximum integer in K_i.

2 Related Work

Social welfare in auctions. When players know exactly their own valuations, the problem of maximizing social welfare in auctions is *optimally* solved (even in multiple good setting) by the well-known *VCG mechanism* [Vic61, Cla71, Gro73]. In this mechanism, reporting the truth is a very-weakly-dominant strategy for every player and, the maximum social welfare is guaranteed at such strategy. In the special case of auctions of a *single* good, the VCG mechanism reduces to the familiar *second-price mechanism.*

Note that the VCG mechanism and, in particular, the second-price mechanism do impose prices on the players, but only as a means to achieve the goal of maximizing social welfare. The alternative goal of *revenue maximization* has been studied extensively, but it is a goal that we do not consider in this work.

Relevant non-Bayesian solution concepts. Since in this paper we focus on non-Bayesian settings, the relevant solution concepts are *dominant* strategies and *undominated* strategies. With regard to dominant strategies, the *revelation principle* (see e.g. [Mye79]) states that if a very-weakly-dominant strategy mechanism exists so does a truthful one (where stating the truth is a very-weakly-dominant strategy). With regard to undominated strategies, Jackson [Jac92, JPS94] points out that studying undominated strategies is meaningful only when considering *bounded* mechanisms. (All undominated strategy mechanisms in our model will be bounded, because they are finite.)

Bayesian models of type uncertainty. Milgrom [Mil89] studies the revenue difference between the second-price and the English auctions in single-good setting, where players do not exactly know their valuations; they know that their valuations are drawn from a common distribution. Porter et al. [PRST08] consider a scheduling problem where tasks are to be scheduled to players, and player i has a private failure rate p_{ij} for task j. This failure rate can be understood as a distribution of player's private type, and this paper studies the dominant strategy mechanisms with respect to the social welfare. Feige and Tennenholtz [FT11] consider another variant of the scheduling problem, where each player i has a task of length l_i to be scheduled and there is only one machine. They study dominant strategy mechanisms *without monetary transfer* in the case that players have uncertainty (and thus a Bayesian distribution) about their own l_i's. They prove that even if each player i has only two possible l_i's, no constant fraction of the maximum social welfare can be guaranteed. To overcome this, they introduce a different measure of social efficiency, which they call "fair share", and provide mechanisms to guarantee an $\Omega(1)$ fair share.[1]

Set-theoretic models of type approximation. The Bayesian model of incomplete information assumes that the players' types are drawn from a joint prior. An alternative way is the set-theoretic approach introduced by Chen and Micali [CM10]: the set of possible types can be considered as a list, and a strategy is dominated only when it is unfavored for all possible types in this list. For instance, this list can be an interval, or can be two different values where the player does not know which is to happen. However, they only model the *external knowledge* (that is a player's knowledge about others') as sets, still assuming that players know their exact true valuations.

In our setting, we use the set-theoretic approach when player do not know their own valuations. We choose to use the terminology "approximate valuations" to emphasize the "Bayesian-free" nature of our model, as opposed to the other natural terminology "uncertain valuations", which carries a Bayesian connotation. (And, of course, our "approximation" has nothing to do with the "approximation factors" considered in, for instance, works that study computing equilibria approximately.)

3 Auctions in the Approximate-Valuation Model

As for any game, an auction can be thought as consisting of two parts, a context and a mechanism. Since our solution concepts are strong enough, the players' beliefs need not be part of our contexts. We only consider auctions of a single good.

Contexts. A *δ-approximate auction context*, where $\delta \in [0, 1)$, consists of the following components.

- $N = \{1, 2, \dots, n\}$, the set of players.
- $\{0, 1, \dots, B\}$, the set of possible valuations for any player. B is the valuation bound.
- tv, the *true-valuation profile*, where each $\mathsf{tv}_i \in \{0, 1, \dots, B\}$.

[1] "ρ fair share" is a property such that each player i has at least ρ success rate if all other players share the same distribution as his l_i.

- $\Omega = (N \cup \{\perp\}) \times \mathbb{R}^N$, the set of outcomes. If $(a, P) \in \Omega$, then $a \in N$ indicates that player a wins the good,and $a = \perp$ that the good remains unallocated; P is the price profile.
- U, the profile of utility functions. Each U_i maps any outcome (a, P) to $\mathsf{tv}_i - P_i$ if $a = i$, and to $-P_i$ otherwise.
- K, the *δ-approximate valuation profile*, where each K_i is a non-empty subset of $[(1-\delta)v_i, (1+\delta)v_i] \cap \{0, 1, \dots, B\}$ for some real number v_i.

Notice that an approximate auction context C is fully specified by n, B, δ, tv, and K: $C = (n, B, \delta, \mathsf{tv}, K)$.

Knowledge. In an approximate auction context $C = (n, B, \delta, \mathsf{tv}, K)$, each player i knows K_i and that $\mathsf{tv}_i \in K_i$, but has no information beyond that. The approximation accuracy δ is typically assumed to be known to the designer. (In a sense, the designer "chooses δ large enough so as to ensure that he is dealing with a δ-approximate context".) However, δ needs not be explicitly known to the players. Indeed, if —say— $K_i = \{90, 91, \dots, 110\}$, then K_i is both a 0.1-approximate and a 0.2-approximate valuation.

Notation. By $\mathscr{C}^{\delta}_{n,B}$ we denote the set of all δ-approximate auction contexts with n players and valuation bound B. For any $x \in \mathbb{R}$, we denote by $\mathsf{int}_\delta(x)$ the "δ-integer-interval centered around x", that is, the set of all integers in $\big[(1-\delta)x, (1+\delta)x\big]$. Finally the *social welfare* (function) *SW* is defined as follows: $\mathrm{SW}(\mathsf{tv}, (a, P)) \stackrel{\text{def}}{=} \mathsf{tv}_a$ for every true-valuation profile tv and outcome (a, P).

Mechanisms. While our contexts have a new component, our mechanisms are totally ordinary. Indeed a mechanism M specifies the players' strategies and a (possibly probabilistic) outcome function. Formally, $M = (\Sigma, F)$ where

- $\Sigma = \Sigma_1 \times \cdots \times \Sigma_n$ where each Σ_i is a finite set, the set of player i's pure strategies; and
- $F \colon \Sigma \to (N \cup \{\perp\}) \times \mathbb{R}^N$ is the *outcome function*, which maps each strategy profile $v \in \Sigma$ to an outcome (a, P) (which may be a random variable if F is probabilistic).

Notation. We denote pure strategies by Latin letters, and possibly mixed strategies by Greek ones. If $s \in \Sigma$ is a pure strategy profile, then we also denote by $F^A_i(s)$ and $F^P_i(s)$ respectively the probability that the good is assigned to player i and the price paid by i under s. We also call F^A the *winning-probability function* of M. If instead $\sigma \in \Delta(\Sigma)$, then $F^A_i(\sigma)$ and $F^P_i(\sigma)$ are the corresponding distributions over winning probabilities and prices for player i. More generally, a winning-probability function is a function $f \colon \Sigma \to [0, 1]^N$ mapping a strategy $s \in \Sigma$ profile to a profile of probabilities, corresponding to the winning probability of each player under strategy s; in particular, $\sum_{i \in N} f_i \le 1$.

4 Dominance in the Approximate-Valuation Model

We need to extend to our approximate valuation setting four classical notions: very weak dominance, weak dominance, dominant strategies, and undominated strategies. The obvious constraint of this extension is that when each approximate valuation K_i consists of a single element, then all extended notions must collapse to the original ones. It should be appreciated that with this single constraint, there are many ways of extending the above notions. Yet, once one keeps "meaningfulness" into account, one really does not have much choice except for the second notion, the weak dominance, where some attention has to be paid in order to maintain the original intuitive meaning that the weakly dominated strategies should not be rationally played (in absence of special beliefs).

Definition 4.1 ([FT91, Definition 1.1])**.** *Fix a player $i \in N$ with approximate valuation K_i. For a (possibly mixed) strategy $\sigma_i \in \Delta(\Sigma_i)$ and a pure strategy $s_i \in \Sigma_i$, we say that*

- σ_i **very-weakly dominates** s_i, *in symbols* $\sigma_i \overset{\text{vw}}{\underset{i,K_i}{\succ}} s_i$, *if*

$$\forall\, \mathsf{tv}_i \in K_i \,,\, \forall\, t_{-i} \in \Sigma_{-i} \,:\, \mathbb{E}U_i(\mathsf{tv}_i, F(\sigma_i \sqcup t_{-i})) \geq \mathbb{E}U_i(\mathsf{tv}_i, F(s_i \sqcup t_{-i})) \;\; .$$

- σ_i **weakly dominates** s_i, *in symbols* $\sigma_i \overset{\text{w}}{\underset{i,K_i}{\succ}} s_i$, *if* σ_i *very-weakly dominates* s_i *and*

$$\exists\, \mathsf{tv}_i \in K_i \,,\, \exists\, t_{-i} \in \Sigma_{-i} \,:\, \mathbb{E}U_i(\mathsf{tv}_i, F(\sigma_i \sqcup t_{-i})) > \mathbb{E}U_i(\mathsf{tv}_i, F(s_i \sqcup t_{-i})) \;\; .$$

For K_i*, the set of* **very-weakly-dominant** *and* **undominated strategies** *are*

$$\mathsf{Dnt}_i(K_i) \stackrel{\text{def}}{=} \left\{ s_i : \forall\, t_i \in \Sigma_i \,,\, s_i \overset{\text{vw}}{\underset{i,K_i}{\succ}} t_i \right\} \quad \textit{and} \quad \mathsf{UDed}_i(K_i) \stackrel{\text{def}}{=} \left\{ s_i : \nexists\, \sigma_i \in \Delta(\Sigma_i) \ \textit{s.t.}\ \sigma_i \overset{\text{w}}{\underset{i,K_i}{\succ}} s_i \right\} \;\; .$$

Finally we set $\begin{cases} \mathsf{Dnt}(K) \stackrel{\text{def}}{=} \mathsf{Dnt}_1(K_1) \times \cdots \times \mathsf{Dnt}_n(K_n) \\ \mathsf{UDed}(K) \stackrel{\text{def}}{=} \mathsf{UDed}_1(K_1) \times \cdots \times \mathsf{UDed}_n(K_n) \end{cases}$.

Remark 4.2.

- Note that, we could have defined "σ_i weakly dominates s_i" using three other quantifications in the additional condition for weak dominance: (1) $\forall \mathsf{tv}_i \forall t_{-i}$, (2) $\exists \mathsf{tv}_i \forall t_{-i}$, or (3) $\forall \mathsf{tv}_i \exists t_{-i}$. None of them, however, is really meaningful: the first two because they do not coincide with the classical notion of weak dominance when K_i is singleton, and the third because it does not capture the "weakest condition" for which a strategy s_i should be discarded in favor of σ_i. To see the latter, notice that we already know that σ_i very-weakly dominates s_i, for player i to discard strategy s_i in favor of σ_i, it should suffice that s_i is strictly worse than σ_i for *even one* of his possible valuation $\mathsf{tv}_i \in K_i$, and we should not insist that this is true for all possible tv_i's.
- We choose to analyze only *very*-weak dominance in *dominant* strategies because, while requiring *weak* dominance is safer, it is unrealistic to hope for such a strong guarantee. (For instance, when players know their valuations exactly, the celebrated second-price mechanism only guarantees that reporting the truth is a *very*-weakly dominant strategy.)
  Furthermore, we choose to analyze only *weak* dominance in *undominated* strategies because, when considering *very-weak* dominance, two equivalent strategies may eliminate each other and the set UDed may become empty.[2]

Let us now state without proof a basic relationship between the above two sets of strategies.

Lemma 4.3. $\mathsf{Dnt}_i(K_i) \neq \emptyset$ *implies* $\mathsf{Dnt}_i(K_i) = \mathsf{UDed}_i(K_i)$.

5 Implementation in the Approximate-Valuation Model

We now extend the basic notions of implementation in dominant strategies and in undominated strategies to the approximate-valuation model. For simplicity, in this paper we shall define them for the problem of maximization of social welfare rather than for a general social choice function.

Definition 5.1. *For every* $\varepsilon \in (0, 1]$*, we define* εMSW *to be the function mapping any true-valuation profile* tv *to* $\varepsilon \cdot \max_{i \in N} \mathsf{tv}_i$.

[2] And, of course, all the mechanism that we will consider are *bounded*, to ensure that implementations in UDed are meaningful [Jac92].

Definition 5.2. *We say that a mechanism $M = (\Sigma, F)$* **partially implements** εMSW **in very-weakly-dominant strategies***, over a class of contexts $\mathscr{C}$, if for all contexts $(n, B, \delta, \mathsf{tv}, K) \in \mathscr{C}$:*

$$\exists\, \sigma \in \mathsf{Dnt}(K)\,, \ \sum_{i \in N} F_i^A(\sigma)\mathsf{tv}_i \geq \varepsilon \mathrm{MSW}(\mathsf{tv}) \ .$$

We say that M **fully implements** εMSW **in undominated strategies***, over a class of contexts $\mathscr{C}$, if for all contexts $(n, B, \delta, \mathsf{tv}, K) \in \mathscr{C}$:*

$$\forall\, \sigma \in \mathsf{UDed}(K)\,, \ \sum_{i \in N} F_i^A(\sigma)\mathsf{tv}_i \geq \varepsilon \mathrm{MSW}(\mathsf{tv}) \ .$$

Remark 5.3.

- We could have also defined *full* implementation of εMSW in very-weakly-dominant strategies, where the social welfare guarantee is required to hold *for all* (rather than for at least one) strategy profile $\sigma \in \mathsf{Dnt}(K)$. However, because $\mathsf{Dnt}(K)$ might be empty, a mechanism M could "vacuously" implement εMSW. On the other hand, if a mechanism M ensures that $\mathsf{Dnt}(K)$ is not empty, then Lemma 4.3 implies that full implementation in very-weakly-dominant strategies becomes equivalent to the notion of full implementation in undominated strategies.
- It would not be meaningful to define *partial* implementation for undominated strategies, because one can trivially engineer a mechanism for which (i) *every* strategy profile consists of undominated strategies, and (ii) every outcome —and including the outcome that maximizes the social welfare— corresponds to at least one strategy profile.

Remark 5.4 (Revelation Principle). We shall rely on the following version of the revelation principle [Mye81], which, as it is easy to verify, continues to hold in the approximate-valuation model. Namely, if a mechanism $M = (\Sigma, F)$ partially implements εMSW in very-weakly-dominant strategies, then there is a "direct" (very-weakly) dominant-strategy truthful mechanism $M' = (\Sigma', F')$ for which Σ' is the set of all δ-approximate valuation profiles, and reporting the true approximate-valuation profile $K \in \Sigma'$ is a very-weakly-dominant strategy profile and guarantees εMSW.

6 Formal Statements of Our Results

Theorem 1. *$\forall n$, $\forall \delta$, and $\forall B > \frac{3-\delta}{2\delta}$, if M partially implements εMSW in very-weakly-dominant strategies over $\mathscr{C}^{\delta}_{n,B}$, then*

$$\varepsilon \leq \frac{1}{n} + \frac{\lfloor \frac{3-\delta}{2\delta} \rfloor + 1}{B} \ .$$

The proof of Theorem 1 can be found in Appendix C.

Theorem 2. *The following two statements hold:*

1. *$\forall n$, $\forall \delta$, and $\forall B$, the second-price mechanism fully implements εMSW in undominated strategies over $\mathscr{C}^{\delta}_{n,B}$ with*

$$\varepsilon \stackrel{\text{def}}{=} \left(\frac{1-\delta}{1+\delta}\right)^2 - \frac{(1-\delta)}{(1+\delta)} \frac{2}{\mathrm{MSW}} \quad \textit{if ties are broken deterministically, and}$$

$$\varepsilon \stackrel{\text{def}}{=} \left(\frac{1-\delta}{1+\delta}\right)^2 \quad \textit{if ties are broken at random (giving positive probability to each tie).}$$

2. *$\forall n$, $\forall \delta$, and $\forall B \geq \frac{1}{\delta}$, if a* **deterministic** *M fully implements εMSW in undominated strategies over $\mathscr{C}^{\delta}_{n,B}$ then*

$$\varepsilon \leq \frac{(1-\delta)^2}{(1+\delta)^2} + \frac{4}{B} \ .$$

The proof of Theorem 2 can be found in Appendix D.

Theorem 3. *The following two statements hold:*

1. *For every n, $\delta \in (0,1)$, and B, there exists a mechanism M_{opt} fully implementing εMSW in undominated strategies over $\mathscr{C}^{\delta}_{n,B}$ with*

$$\varepsilon \stackrel{\text{def}}{=} \frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} \ .$$

2. *For every n, $\delta \in (0,1)$, and $B \geq \frac{1}{\delta}$, if M is a (deterministic or probabilistic) mechanism fully implementing εMSW in undominated strategies over $\mathscr{C}^{\delta}_{n,B}$ then*

$$\varepsilon \leq \frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} + \frac{4}{B} \ .$$

The proof of Theorem 3 can be found in Appendix E. This theorem actually is the technically hardest of our results. However, our optimal mechanism M_{opt} is quite simple to play. For convenience, its concise description is given in Appendix A.

References

[Cla71] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971. 5

[CLS+10] Wei Chen, Pinyan Lu, Xiaorui Sun, Yajun Wang, and Zeyuan Allen Zhu. Pricing in social networks: Equilibrium and revenue maximization. *ArXiv e-prints*, abs/1007.1501, March 2010. 22

[CM10] Jing Chen and Silvio Micali. The conservative model of incomplete information and the second-knowledge mechanism. Technical report, MIT, October 2010. 1, 6

[FT91] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991. 7

[FT11] Uriel Feige and Moshe Tennenholtz. Mechanism design with uncertain inputs (to err is human, to forgive divine). *ArXiv e-prints*, March 2011. arXiv:1103.2520v1. 6

[Gro73] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973. 5

[Jac92] Matthew O Jackson. Implementation in undominated.strategies: A look at bounded mechanisms. *Review of Economic Studies*, 59(4):757–75, October 1992. 6, 8, 13

[JPS94] Matthew O. Jackson, Thomas Palfrey, and Sanjay Srivastava. Undominated Nash implementation in bounded mechanisms. *Games and Economic Behavior*, 6(3):474–501, 1994. 6

[Mil89] Paul Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22, Summer 1989. 6

[Mye79] Roger B Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, 47(1):61–73, January 1979. 6

[Mye81] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. 9

[PRST08] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172:1783–1799, October 2008. 6

[Vic61] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961. 5

Appendix

A The Optimal Mechanism M_{opt}

Consider an approximate auction of a single good with n players and valuation bound B.

A.1 The Strategy Sets of M_{opt}

For every player $i \in N$, we set $\Sigma_i \stackrel{\text{def}}{=} \{0, \ldots, B\}$.

A.2 The Outcome Function of M_{opt}

First consider the following winning-probability function:

Definition A.1. *For every $\delta \in (0,1)$, and let $D_\delta \stackrel{\text{def}}{=} \left(\frac{1+\delta}{1-\delta}\right)^2 - 1 > 0$. We define the function $f^{(\delta)} \colon [0,B]^N \to [0,1]^N$ as follows:*

- *for every $z = (z_1, \ldots, z_n) \in [0,B]^N$ such that $z_1 \geq z_2 \geq \cdots \geq z_n$, let n^* be the least index in N such that*

$$\forall\, i \in \{1, \ldots, n^*\}, \quad z_i \leq \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta} \ ,$$

  call players $1, \ldots, n^*$ the winners *and players $n^*+1, \ldots, n$ the* losers*, and then set*

$$f_i^{(\delta)}(z) \stackrel{\text{def}}{=} \begin{cases} \frac{1}{n} \cdot \frac{n+D_\delta}{n^*+D_\delta} \cdot \frac{z_i(n^*+D_\delta) - \sum_{j=1}^{n^*} z_j}{z_i D_\delta}, & \text{if } i \leq n^*, \\ 0, & \text{if } i > n^*; \end{cases}$$

- *for other z, define $f^{(\delta)}$ by extending it symmetrically.*

The (code for the) outcome function $F_{f^{(\delta)}}$ of our mechanism $M_{\mathsf{opt}} = (\{0,1,\ldots,B\}^N, F_{f^{(\delta)}})$ is:

Code for outcome function of M_{opt}

public parameter: $\delta \in (0,1)$

inputs: $v_1, \ldots, v_n \in \{0, 1, \ldots, B\}$

output: (i, P), where $i \in N \cup \{\perp\}$ is the winning player and $P \in \mathbb{R}^N$ is the price profile

pseudocode:

$F_{f^{(\delta)}}(v_1, \ldots, v_n) \stackrel{\text{def}}{=}$

1. Draw r uniformly at random in $[0,1]$.
2. Define $f_0^{(\delta)} \stackrel{\text{def}}{=} 0$.
3. If there exists $i \in N$ such that $\sum_{j=0}^{i-1} f_i^{(\delta)}(v) < r \leq \sum_{j=0}^{i} f_i^{(\delta)}(v)$:
   - Compute $P_i \stackrel{\text{def}}{=} v_i - \frac{\int_0^{v_i} f_i^{(\delta)}(z \sqcup v_{-i})\, dz}{f_i^{(\delta)}(v)}$, and $P_j = 0$ for $j \neq i$, and output (i, P).
4. Otherwise, output $(\perp, (0, \ldots, 0))$. (No player is assigned the good.)

We note that our mechanism can be tweaked to make sure that the good is always assigned to some player. But the proof is more involved than it already is, and we leave it to a future version of this paper.

B Two Structural Lemmas

In this section we establish two basic lemmas. The first crucial to our "negative" results, the second crucial to our "positive" ones.

B.1 The Undominated Intersection Lemma

To prove that a given social choice function cannot be implemented in undominated strategies in the approximate-valuation model, we need to establish some basic structural properties about undominated strategies. If two possible approximate-valuation sets of the same player i, K_i and $\widetilde{K}_i$, intersected and any possible valuation were a strategy available to i, then there would be a simple relationship between their corresponding seta of *pure* undominated strategies: namely, $K_i \cap \widetilde{K}_i \neq \emptyset \rightarrow \mathsf{UDed}_i(K_i) \cap \mathsf{UDed}_i(\widetilde{K}_i) \neq \emptyset$. However, a mechanism can choose the strategies available to each player in a totally arbitrary manner, and thus such an attractively simple relationship may not hold at all: as soon as K_i and $\widetilde{K}_i$ are even slightly different, their corresponding sets of undominated strategies may be totally unrelated. We prove however that, no matter what the available strategies may be, whenever K_i and $\widetilde{K}_i$ have at least two values in common, there exist two "almost payoff-equivalent" strategies σ_i and $\widetilde{\sigma}_i$ respectively in $\Delta(\mathsf{UDed}_i(K_i))$ and $\Delta(\mathsf{UDed}_i(\widetilde{K}_i))$. This relationship will be sufficient to prove all of our impossibility results for implementation in undominated strategies.

Lemma B.1 (Undominated Intersection Lemma)**.** *Fix a mechanism M. For every player $i \in N$, if K_i and $\widetilde{K}_i$ intersect in at least two points, then, for every $\varepsilon > 0$, there exists some mixed strategy $\sigma_i \in \Delta(\mathsf{UDed}_i(K_i))$ and $\widetilde{\sigma}_i \in \Delta(\mathsf{UDed}_i(\widetilde{K}_i))$ such that*

$$\forall\, s_{-i} \in \Sigma_{-i}\,, \quad \begin{aligned} &\left|F_i^A(\sigma_i \sqcup s_{-i}) - F_i^A(\widetilde{\sigma}_i \sqcup s_{-i})\right| < \varepsilon \\ &\left|F_i^P(\sigma_i \sqcup s_{-i}) - F_i^P(\widetilde{\sigma}_i \sqcup s_{-i})\right| < \varepsilon \end{aligned}$$

Proof. Let i be a player; let K_i and $\widetilde{K}_i$ be two approximate-valuation sets of i; let x_i and y_i be two distinct elements in $K_i \cap \widetilde{K}_i$; and, without loss of generality, let $x_i > y_i$. Note that both $\mathsf{UDed}_i(K_i)$ and $\mathsf{UDed}_i(\widetilde{K}_i)$ are nonempty.

If there exists a common (pure) strategy $s_i \in \mathsf{UDed}_i(K_i) \cap \mathsf{UDed}_i(\widetilde{K}_i)$, then we can simply set $\sigma_i \stackrel{\text{def}}{=} \widetilde{\sigma}_i \stackrel{\text{def}}{=} s_i$ and we are done. Therefore, we now consider the case where $\mathsf{UDed}_i(K_i)$ and $UWD_i(\widetilde{K}_i)$ are disjoint.

Let s_i be any strategy in $\mathsf{UDed}_i(K_i)$, so that s_i is *not* in $\mathsf{UDed}_i(\widetilde{K}_i)$ and thus, by definition, we know that $\widetilde{\sigma}_i \underset{i,\widetilde{K}_i}{\overset{\text{w}}{\succ}} s_i$ for some $\widetilde{\sigma}_i \in \Delta(\mathsf{UDed}_i(\widetilde{K}_i))$.[3] We now argue that

$$\exists\, \tau_i \in \Delta(\mathsf{UDed}_i(K_i)) \text{ such that } \tau_i \underset{i,K_i}{\overset{\text{w}}{\succ}} \widetilde{\sigma}_i \ . \tag{B.1}$$

Note that, while we have only defined what it means for a *pure* strategy to be dominated by a possibly mixed one, the definition trivially extends to the case of dominated strategies that are *mixed*, as is the case in "$\tau_i \underset{i,K_i}{\overset{\text{w}}{\succ}} \widetilde{\sigma}_i$" in Equation B.1. Writing $\widetilde{\sigma}_i = \sum_j \alpha^{(j)} s_i^{(j)}$ (where the summation is over a finite set) and invoking again the disjointness of the two undominated strategy sets, we

[3] This is because in our definition a mechanism M can only provide finite number of strategies. This implies that M is bounded (see [Jac92]), i.e., if a strategy is weakly dominated, it is also weakly dominated by some mixed strategy within $\Delta(\mathsf{UDed}_i)$.

obtain, for each j, a $\tau_i^{(j)} \in \Delta(\mathsf{UDed}_i(K_i))$ such that $\tau_i^{(j)} \underset{i,K_i}{\overset{\mathrm{w}}{\succ}} \widetilde{s}_i^{(j)}$; we can then set $\tau = \sum_j \alpha^{(j)} \tau_i^{(j)}$ and deduce that $\tau_i \underset{i,K_i}{\overset{\mathrm{w}}{\succ}} \widetilde{\sigma}_i$.

For the same reason, we can also find some $\widetilde{\tau}_i \in \Delta(\mathsf{UDed}_i(\widetilde{K}_i))$ such that $\widetilde{\tau}_i \underset{i,\widetilde{K}_i}{\overset{\mathrm{w}}{\succ}} \tau_i$. Continuing in this fashion, "jumping" back and forth between $\Delta(\mathsf{UDed}_i(K_i))$ and $\Delta(\mathsf{UDed}_i(\widetilde{K}_i))$, we obtain an infinite chain of (possibly recurring) strategies. So consider such a chain $\{\sigma_i^{(k)}, \widetilde{\sigma}_i^{(k)}\}_{k\in\mathbb{N}}$; by construction,

$$\dots \underset{i,\widetilde{K}_i}{\overset{\mathrm{w}}{\succ}} \widetilde{\sigma}_i^{(2)} \underset{i,K_i}{\overset{\mathrm{w}}{\succ}} \sigma_i^{(2)} \underset{i,\widetilde{K}_i}{\overset{\mathrm{w}}{\succ}} \widetilde{\sigma}_i^{(1)} \underset{i,K_i}{\overset{\mathrm{w}}{\succ}} \sigma_i^{(1)} \ .$$

By the definition of dominance, we know that, for all $s_{-i} \in \Sigma_{-i}$:

$$\begin{aligned} \forall\, k \in \mathbb{N} \quad \forall \widetilde{\mathsf{tv}}_i \in \widetilde{K}_i,\quad & F_i^A(\sigma_i^{(k)})\widetilde{\mathsf{tv}}_i - F_i^P(\sigma_i^{(k)}) \;\le\; F_i^A(\widetilde{\sigma}_i^{(k)})\widetilde{\mathsf{tv}}_i - F_i^P(\widetilde{\sigma}_i^{(k)}) \\ \forall \mathsf{tv}_i \in K_i,\quad & F_i^A(\widetilde{\sigma}_i^{(k)})\mathsf{tv}_i - F_i^P(\widetilde{\sigma}_i^{(k)}) \;\le\; F_i^A(\sigma_i^{(k+1)})\mathsf{tv}_i - F_i^P(\sigma_i^{(k+1)}) \end{aligned} \tag{B.2}$$

where for instance $F_i^A(\sigma_i^{(k)})$ is an abbreviation for $F_i^A(\sigma_i^{(k)} \sqcup s_{-i})$.

Because we have $\{x_i, y_i\} \subset K_i \cap \widetilde{K}_i$, we can set both tv_i and $\widetilde{\mathsf{tv}}_i$ to be x_i for all $k \in \mathbb{N}$. Fixing any s_{-i}, Equation B.2 can be re-written as follows:

$$\begin{aligned} \forall\, k \in \mathbb{N} \quad F_i^A(\sigma_i^{(k)})x_i - F_i^P(\sigma_i^{(k)}) &\le F_i^A(\widetilde{\sigma}_i^{(k)})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}) \\ F_i^A(\widetilde{\sigma}_i^{(k)})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}) &\le F_i^A(\sigma_i^{(k+1)})x_i - F_i^P(\sigma_i^{(k+1)}) \end{aligned}$$

This gives us an infinite sequence of non-decreasing numbers. Notice that the sequence is also upper bounded because the winning-probability function F^A has a value between 0 and 1, and the price is non-negative; hence each term is upper bounded by B. By the Bolzano-Weierstrass theorem, there must exist some $H_\varepsilon \in \mathbb{N}$ such that

$$\forall\, k > H_\varepsilon \qquad F_i^A(\sigma_i^{(k)})x_i - F_i^P(\sigma_i^{(k)}) \le F_i^A(\widetilde{\sigma}_i^{(k)})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}) \tag{B.3}$$

$$F_i^A(\widetilde{\sigma}_i^{(k)})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}) \le F_i^A(\sigma_i^{(k+1)})x_i - F_i^P(\sigma_i^{(k+1)}) \tag{B.4}$$

$$F_i^A(\sigma_i^{(k+1)})x_i - F_i^P(\sigma_i^{(k+1)}) \le F_i^A(\sigma_i^{(k)})x_i - F_i^P(\sigma_i^{(k)}) + \frac{\varepsilon}{3B} \tag{B.5}$$

While this H_ε is defined for a specific choice of $s_{-i} \in \Sigma_{-i}$, we can take the largest H_ε among all s_{-i} because Σ_{-i} is finite; this makes Equation B.3, Equation B.4 and Equation B.5 hold for all s_{-i}.

Similarly, we can choose both tv_i and $\widetilde{\mathsf{tv}}_i$ to be y_i, obtaining another H_ε (and thus we take the larger one of the two), so that for all s_{-i} we have not only Equation B.3, Equation B.4 and Equation B.5 but also

$$\forall k > H_\varepsilon \qquad F_i^A(\sigma_i^{(k)})y_i - F_i^P(\sigma_i^{(k)}) \le F_i^A(\widetilde{\sigma}_i^{(k)})y_i - F_i^P(\widetilde{\sigma}_i^{(k)}) \tag{B.6}$$

$$F_i^A(\widetilde{\sigma}_i^{(k)})y_i - F_i^P(\widetilde{\sigma}_i^{(k)}) \le F_i^A(\sigma_i^{(k+1)})y_i - F_i^P(\sigma_i^{(k+1)}) \tag{B.7}$$

$$F_i^A(\sigma_i^{(k+1)})y_i - F_i^P(\sigma_i^{(k+1)}) \le F_i^A(\sigma_i^{(k)})y_i - F_i^P(\sigma_i^{(k)}) + \frac{\varepsilon}{3B} \tag{B.8}$$

Next, we pick an arbitrary $k > H_\varepsilon$, and will prove that the two strategies $\sigma_i^{(k+1)}$ and $\widetilde{\sigma}_i^{(k)}$ are the two strategies that we are looking for.

We first sum up Equation B.3, Equation B.5 and Equation B.7. The (expected) prices and the $F_i^A(\sigma_i^{(k)})x_i$ term will cancel and we can deduce that

$$F_i^A(\sigma_i^{(k+1)})(x_i - y_i) \le F_i^A(\widetilde{\sigma}_i^{(k)})(x_i - y_i) + \frac{\varepsilon}{3B} \ .$$

We then sum up Equation B.4, Equation B.6 and Equation B.8. The (expected) prices and the $F_i^A(\sigma_i^{(k)})y_i$ term will cancel and we can deduce that

$$F_i^A(\widetilde{\sigma}_i^{(k)})(x_i - y_i) \le F_i^A(\sigma_i^{(k+1)})(x_i - y_i) + \frac{\varepsilon}{3B} \ .$$

Since $x_i - y_i \ge 1$, combining the last two inequalities yields

$$\left|F_i^A(\widetilde{\sigma}_i^{(k)}) - F_i^A(\sigma_i^{(k+1)})\right| \le \frac{\varepsilon}{3B} \ . \tag{B.9}$$

We emphasize again that this holds for all s_{-i} and our notion $F_i^A(\widetilde{\sigma}_i^{(k)})$ is an abbreviation to $F_i^A(\widetilde{\sigma}_i^{(k)} \sqcup s_{-i})$, so the first inequality in Lemma B.1 is derived.

We now consider the price terms. Combining Equation B.4 and Equation B.9, we get:

$$\begin{aligned} F_i^A(\widetilde{\sigma}_i^{(k)})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}) &\le F_i^A(\sigma_i^{(k+1)})x_i - F_i^P(\sigma_i^{(k+1)}) \\ &\le (F_i^A(\widetilde{\sigma}_i^{(k)}) + \frac{\varepsilon}{3B})x_i - F_i^P(\sigma_i^{(k+1)}) \\ &\Rightarrow -F_i^P(\widetilde{\sigma}_i^{(k)}) \le \frac{\varepsilon}{3} - F_i^P(\sigma_i^{(k+1)}) \ . \end{aligned} \tag{B.10}$$

Summing up Equation B.3 and Equation B.5 and then substituting Equation B.9, we get:

$$\begin{aligned} F_i^A(\sigma_i^{(k+1)})x_i - F_i^P(\sigma_i^{(k+1)}) &\le F_i^A(\widetilde{\sigma}_i^{(k)})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}) + \frac{\varepsilon}{3B} \\ &\le (F_i^A(\sigma_i^{(k+1)}) + \frac{\varepsilon}{3B})x_i - F_i^P(\widetilde{\sigma}_i^{(k)}) + \frac{\varepsilon}{3B} \\ &\Rightarrow -F_i^P(\sigma_i^{(k+1)}) \le \frac{2\varepsilon}{3} - F_i^P(\widetilde{\sigma}_i^{(k)}) \ . \end{aligned} \tag{B.11}$$

Combining the last two inequalities Equation B.10 and Equation B.11 we immediately get the second requirement of Lemma B.1, since this holds for *all* s_{-i}:

$$\left|F_i^P(\widetilde{\sigma}_i^{(k)}) - F_i^P(\sigma_i^{(k+1)})\right| \le \frac{2\varepsilon}{3} \ ,$$

therefore obtaining the second inequality claimed by the lemma, and thus completing the proof of the lemma. □

B.2 The Distinguishable Monotonicity Lemma

To prove that a given social choice function can be implemented in undominated strategies, we are happy to work with mechanisms using very special strategies and allocation functions. The important thing is that they ensure that the undominated strategies corresponding to a given approximate-valuation set are simple to characterize. Specifically, we consider mechanisms in which any possible valuation is an available strategy, and whose winning-probability functions are restrictions (to $\{0, \ldots, B\}^N$) of integrable functions (over $[0, B]^N$) satisfying a suitable monotonicity property. A simple lemma will then guarantee that the undominated strategies of such mechanisms are easy to work with. (We stress that our setting continues to be discrete, as the analysis over a continuous domain is only a tool for proving the lemma.)

Definition B.2. *Let $f : [0,B]^N \to [0,1]^N$ be a winning-probability function. We call a mechanism $M = (\Sigma, F)$ an **f-mechanism** if $\Sigma = \{0,1,\dots,B\}^N$ and:*

- *f is integrable and monotonic (i.e., each f_i is integrable and a non-decreasing function in player i's coordinate), and*
- *$F^A = f|_{\{0,1,\dots,B\}}$ and, for all $i \in N$ and $v \in \Sigma$, $F_i^P(v) = v_i \cdot f_i(v_i \sqcup v_{-i}) - \int_0^{v_i} f_i(z \sqcup v_{-i})\, dz$.*

Notation. For a given winning-probability function f, we will denote by M_f the corresponding f-mechanism. Note that M_f is deterministic if and only if $f(\Sigma) = F^A(\Sigma) \subseteq \{0,1\}^N$.

In the definition of an f-mechanism we have only required the *expected* price imposed by the mechanism to satisfy a certain condition depending on f. The underlying (possibly probabilistic) price function will not matter for our analysis (and it can be easily chosen to ensure other desirable properties of the f-mechanism, such as the fact that the opt-out condition is satisfied).

The lemma below shows that, if the winning-probability function of M_f, beyond being non-decreasing, actually satisfies an additional monotonicity condition, then we can simply characterize the undominated strategies of M_f associated to a given approximate-valuation sets, so as to enable us to engineer implementations in undominated strategies.

Definition B.3. *Let $f : [0,B]^N \to [0,1]^N$ be a winning-probability function. For $d \in \{1,2\}$, we say that f is **d-distinguishably monotonic** (d-DM, for short) if f is integrable and monotonic, and, in addition, it satisfies the following "distinguishability" condition:*

$$\forall i \in N,\ \forall v_i, v_i' \in \Sigma_i \text{ s.t. } v_i \le v_i' - d,\ \exists\, v_{-i} \in \Sigma_{-i} \quad \int_{v_i}^{v_i'} \big(f_i(z \sqcup v_{-i}) - f_i(v_i \sqcup v_{-i})\big)\, dz > 0 \ .$$

Distinguishable monotonicity is certainly a requirement additional to the traditional one of non-decreasing monotonicity, but is actually quite mild. (Indeed, the second-price mechanism can be represented as an f-mechanism where f is 2-DM —and if ties are broken at random, it is even 1-DM.) Nonetheless, it is a quite useful requirement.

Lemma B.4 (Distinguishable Monotonicity Lemma)**.** *For all f-mechanisms M_f, players i, and δ-approximate-valuation profiles K,*

$$\begin{aligned} &\mathsf{UDed}_i(K_i) \subseteq \{\min K_i, \dots, \max K_i\} \ \text{ if } f \text{ is 1-DM, and} \\ &\mathsf{UDed}_i(K_i) \subseteq \{\min K_i - 1, \dots, \max K_i + 1\} \ \text{ if } f \text{ is 2-DM.} \end{aligned}$$

(Above $\min K_i$ and $\max K_i$ respectively denote the minimum and maximum integers in K_i.)

Proof. Suppose f is d-DM, for $d \in \{1,2\}$, and define the strategy $v_i^* \stackrel{\text{def}}{=} \min K_i$. Then,

- *v_i^* very-weakly dominates every $v_i \le v_i^* - d$.* Fix any (pure) strategy sup-profile $v_{-i} \in \Sigma_{-i}$ for the other players and any possible true valuation $\mathsf{tv}_i \in K_i$. Letting $v^* = v_i^* \sqcup v_{-i}$ and $v = v_i \sqcup v_{-i}$, we prove that

$$\begin{aligned} &\mathbb{E}U_i\big(\mathsf{tv}_i, F(v^*)\big) - \mathbb{E}U_i\big(\mathsf{tv}_i, F(v)\big) \\ &= \big(f_i(v^*) - f_i(v)\big) \cdot \mathsf{tv}_i - \big(F_i^P(v^*) - F_i^P(v)\big) \\ &= \big(f_i(v^*) - f_i(v)\big) \cdot \mathsf{tv}_i - \left(v_i^* \cdot f_i(v^*) - \int_0^{v_i^*} f_i(z \sqcup v_{-i})\, dz - v_i \cdot f_i(v) + \int_0^{v_i} f_i(z \sqcup v_{-i})\, dz\right) \\ &= \big(f_i(v^*) - f_i(v)\big) \cdot (\mathsf{tv}_i - v_i^*) + \int_{v_i}^{v_i^*} \big(f_i(z \sqcup v_{-i}) - f_i(v)\big)\, dz \ . \end{aligned}$$

Now note that, since $\mathsf{tv}_i \in K_i$, $\mathsf{tv}_i - v_i^* = \mathsf{tv}_i - \min K_i \geq 0$; moreover, by the monotonicity of f, whenever $z \geq v_i$, it holds that $f_i(z \sqcup v_{-i}) \geq f_i(v)$. We deduce that $\mathbb{E}U_i\big(\mathsf{tv}_i, F(v^*)\big) \geq \mathbb{E}U_i\big(\mathsf{tv}_i, F(v)\big)$. We conclude that v_i^* very-weakly dominates v_i.

- *There is a strategy sub-profile v_{-i} for which v_i^* is strictly better than every $v_i \leq v_i^* - d$.* Indeed, since f is d-DM, for every $v_i \leq v_i^* - d$ there exists a strategy sub-profile v_{-i} to make $\int_{v_i}^{v_i^*} \big(f_i(z \sqcup v_{-i}) - f_i(v)\big)\, dz$ strictly positive, so for this v_{-i} playing v_i^* is strictly better than v_i.

An analogous argument shows that the strategy $v_i^* = \max K_i$ weakly dominates every strategy $v_i \in \Sigma_i$ with $v_i \geq v_i^* + d$. Therefore, the (pure) undominated strategies must lie in the interval $\{\min K_i - (d-1), \ldots, \max K_i + (d-1)\}$. This completes the proof of the theorem. $\square$

C Proof for Theorem 1

Fix arbitrarily a number of players n, an approximation accuracy δ, and then a valuation bound $B > \frac{3-\delta}{2\delta}$. In light of the Revelation Principle in the approximate-valuation model (see Remark 5.4), it suffices to prove that

For every (possibly probabilistic) direct dominant strategy truthful mechanism $M = (\Sigma, F)$ over $\mathscr{C}^{\delta}_{n,B}$, there exists a δ-approximate-valuation profile K and a true-valuation profile $\mathsf{tv} \in K$ such that

$$\mathbb{E}\big[\mathrm{SW}\big(\mathsf{tv}, F(K)\big)\big] \leq \left(\frac{1}{n} + \frac{\lfloor (3-\delta)/2\delta \rfloor + 1}{B}\right) \mathrm{MSW}(\mathsf{tv}) \ .$$

We start by proving a separate claim: essentially, if a player reports a δ-integer-interval whose center is sufficiently high, then his winning probability and expected price remain constant.

Claim C.1. *For all players i, integers $x \in (\frac{3-\delta}{2\delta}, B]$, and δ-approximate-valuation sub-profiles K_{-i},*

$$F_i^A(\mathsf{int}_\delta(x) \sqcup K_{-i}) = F_i^A(\mathsf{int}_\delta(x+1) \sqcup K_{-i}) \textit{ and}$$
$$F_i^P(\mathsf{int}_\delta(x) \sqcup K_{-i}) = F_i^P(\mathsf{int}_\delta(x+1) \sqcup K_{-i}) \ .$$

Proof. Define $K_i \stackrel{\text{def}}{=} \mathsf{int}_\delta(x)$ and $K_i' \stackrel{\text{def}}{=} \mathsf{int}_\delta(x+1)$. Then:

- If player i has approximate-valuation set K_i then reporting K_i very-weakly dominates reporting K_i':

$$\begin{aligned} \forall\, \mathsf{tv}_i \in K_i : \quad & F_i^A(K_i \sqcup K_{-i}) \cdot \mathsf{tv}_i - F_i^P(K_i \sqcup K_{-i}) \\ & \geq F_i^A(K_i' \sqcup K_{-i}) \cdot \mathsf{tv}_i - F_i^P(K_i' \sqcup K_{-i}) \ . \end{aligned}$$

- If player i has approximate-valuation set K_i' then reporting K_i' very-weakly dominates reporting K_i:

$$\begin{aligned} \forall\, \mathsf{tv}_i' \in K_i' : \quad & F_i^A(K_i' \sqcup K_{-i}) \cdot \mathsf{tv}_i' - F_i^P(K_i' \sqcup K_{-i}) \\ & \geq F_i^A(K_i \sqcup K_{-i}) \cdot \mathsf{tv}_i' - F_i^P(K_i \sqcup K_{-i}) \ . \end{aligned}$$

On one hand, we can choose $\mathsf{tv}_i = x$ and $\mathsf{tv}'_i = x+1$, and sum the two corresponding inequalities. The F_i^P price terms cancel, and we get:

$$F_i^A(K'_i \sqcup K_{-i}) \geq F_i^A(K_i \sqcup K_{-i}) \enspace .$$

On the other hand, we can choose $\mathsf{tv}_i = \lfloor x(1+\delta) \rfloor$ and $\mathsf{tv}'_i = \lceil (x+1)(1-\delta) \rceil$,[4] and sum the two corresponding inequalities to get:

$$\big(F_i^A(K_i \sqcup K_{-i}) - F_i^A(K'_i \sqcup K_{-i})\big) \cdot \big(\lfloor x(1+\delta) \rfloor - \lceil (x+1)(1-\delta) \rceil\big) \geq 0 \enspace .$$

Therefore, whenever $x > \frac{3-\delta}{2\delta}$, we always have $\lfloor x(1+\delta) \rfloor - \lceil (x+1)(1-\delta) \rceil > 0$ and thus $F_i^A(K_i \sqcup K_{-i}) = F_i^A(K'_i \sqcup K_{-i})$. Finally, going back to the two inequalities for very weak dominance, we also deduce that $F_i^P(K_i \sqcup K_{-i}) = F_i^P(K'_i \sqcup K_{-i})$, as desired. $\square$

We can now go back to the proof of Theorem 1.

Choose the profile of approximate-valuation sets $K \stackrel{\text{def}}{=} (\mathsf{int}_\delta(c), \mathsf{int}_\delta(c), \dots, \mathsf{int}_\delta(c))$, where $c \stackrel{\text{def}}{=} \lfloor \frac{3-\delta}{2\delta} \rfloor + 1$. By averaging, because the summation of $F_i^A(K)$ over $i \in N$ cannot be greater than 1, there must exist a player j such that $F_j^A(K) \leq 1/n$. Without loss of generality, let such player be player 1. Then, invoking Claim C.1 multiple times with player $i = 1$, $K_{-i} = (\mathsf{int}_\delta(c), \mathsf{int}_\delta(c), \dots, \mathsf{int}_\delta(c))$ of this proof, and x being $c, c+1, \dots, B$, we obtain that

$$F_1^A(\mathsf{int}_\delta(B), \mathsf{int}_\delta(c), \dots, \mathsf{int}_\delta(c)) = F_1^A(K) \leq \frac{1}{n} \enspace .$$

Now suppose that the true approximate-valuation profile of the players is $K' \stackrel{\text{def}}{=} (\mathsf{int}_\delta(B), \mathsf{int}_\delta(c), \dots, \mathsf{int}_\delta(c))$. Then, for the choice of true-valuation profile $\mathsf{tv} = (B, c, \dots, c) \in K'$, we get the following social welfare:

$$\mathbb{E}\big[\mathrm{SW}\big(\mathsf{tv}, F(K')\big)\big] \leq \frac{1}{n}B + \frac{n-1}{n}c \leq \left(\frac{1}{n} + \frac{c}{B}\right) B = \left(\frac{1}{n} + \frac{c}{B}\right) \cdot \mathrm{MSW}(\mathsf{tv}) \enspace ,$$

as desired.

D Proof for Theorem 2

D.1 Proof of Statement 1

Let us first consider the second-price mechanism with a *deterministic* tie-breaking rule, $M_{\mathsf{2P}} = (\Sigma_{\mathsf{2P}}, F_{\mathsf{2P}})$, and prove that

For every n, δ, B. For every δ-approximate-valuation profile $K = (K_1, \dots, K_n)$, every $\mathsf{tv} \in K$, and every $v \in \mathsf{UDed}(\mathsf{tv})$, the following inequality holds:

$$\mathrm{SW}\big(\mathsf{tv}, F_{\mathsf{2P}}(v)\big) \geq \frac{(1-\delta)^2}{(1+\delta)^2}\mathrm{MSW}(\mathsf{tv}) - 2 \cdot \frac{1-\delta}{1+\delta} \enspace . \tag{D.1}$$

By the definition of δ-approximate-valuation set, for each player i, we can assume that $K_i \subset [x_i(1-\delta), x_i(1+\delta)] \cap \mathbb{Z}$ for some x_i.

[4] As long as $x > \frac{1}{2\delta}$ it is guaranteed that, for these choices, $\mathsf{tv}_i \in K_i$ and $\mathsf{tv}'_i \in K'_i$. But later we will choose $x > \frac{3-\delta}{2\delta}$, so we are safe.

We observe that for each player i, the set of undominated strategies

$$\mathsf{UDed}_i(x) \subset \{\min K_i - 1, \ldots, \max K_i + 1\} \subset \big[\lceil x_i(1-\delta)\rceil - 1, \lfloor x_i(1+\delta)\rfloor + 1\big] \cap \mathbb{Z} \ .$$

The first inclusion follows from the Distinguishable Monotonicity Lemma (Lemma B.4), combining the fact that the second-price mechanism M_{2P} is an f-mechanism where f is 2-DM (recall the explanation after Definition B.3).

Now we prove the lower bound of the social welfare. Consider a possible true-valuation profile $\mathsf{tv} \in K$, in which player i^* has the highest valuation $\mathsf{tv}_{i^*} = \max_i \mathsf{tv}_i$. For an arbitrary possible (pure) strategy profile $v \in \mathsf{UDed}(K)$, and consider the winner j^* with the highest bid $v_{j^*} = \max_j v_j$. We now bound the difference between tv_{i^*} and tv_{j^*}. We only need to consider the case when $i^* \neq j^*$.

Our observation suggests that $v_{i^*} \geq \lceil x_{i^*}(1-\delta)\rceil - 1$, and $v_{j^*} \leq \lfloor x_{j^*}(1+\delta)\rfloor + 1$. The fact that j^* is the winner implies $v_{j^*} \geq v_{i^*}$. Combining them we have $x_{i^*}(1-\delta) \leq x_{j^*}(1+\delta) + 2$. Since we also know that $\mathsf{tv}_{i^*} \leq x_{i^*}(1+\delta)$ and $\mathsf{tv}_{j^*} \geq x_{j^*}(1-\delta)$, all of the above suggests that:

$$\begin{aligned}
\mathrm{SW}(\mathsf{tv}, F_{\mathsf{2P}}(v)) = \mathsf{tv}_{j^*} &\geq (1-\delta)x_{j^*} \geq (1-\delta)\frac{1-\delta}{1+\delta}x_{i^*} - \frac{2(1-\delta)}{(1+\delta)} \\
&\geq (1-\delta)\frac{1-\delta}{1+\delta}\frac{1}{1+\delta}\mathsf{tv}_{i^*} - \frac{2(1-\delta)}{(1+\delta)} = \frac{(1-\delta)^2}{(1+\delta)^2}\mathrm{MSW}(\mathsf{tv}) - \frac{2(1-\delta)}{(1+\delta)} \ .
\end{aligned}$$

Thus, Statement 1 of Theorem 2 holds if second-price mechanism breaks ties deterministically.

Consider now the case where the second-price mechanism breaks ties at random (assigning a positive probability to each tie). Then, one can use a proof analogous to the one above, with the only difference being that M_{2P} is an f-mechanism where f is 1-DM (instead of only 2-DM), and invoking a stronger result of Lemma B.4, to show the stronger lower bound

$$\mathrm{SW}\big(\mathsf{tv}, F_{\mathsf{2P}}(v)\big) \geq \frac{(1-\delta)^2}{(1+\delta)^2}\mathrm{MSW}(\mathsf{tv}) \ .$$

Thus Statement 1 is true in all cases. ■

D.2 Proof of Statement 2

Fix arbitrarily a number of players n, an approximation accuracy $\delta \in (0,1)$, and then a valuation bound $B \geq \frac{1}{\delta}$. We need to prove that:

For every deterministic *mechanism $M = (\Sigma, F)$ over $\mathscr{C}^{\delta}_{n,B}$, there exists a δ-approximate-valuation profile K, a strategy profile $s \in \mathsf{UDed}(K)$ and a true-valuation profile* $\mathsf{tv} \in K$ *such that*

$$\mathrm{SW}\big(\mathsf{tv}, F(s)\big) \leq \left(\frac{(1-\delta)^2}{(1+\delta)^2} + \frac{4}{B}\right)\mathrm{MSW}(\mathsf{tv}) \ .$$

Define $x \stackrel{\text{def}}{=} B$ and $y \stackrel{\text{def}}{=} \lfloor\frac{x(1-\delta)+2}{1+\delta}\rfloor$. By our choice of B, we know that $x \geq y$. Therefore both $x, y \in \{0, 1, \ldots, B\}$. Furthermore, one can verify that $\mathsf{int}_\delta(x) \cap \mathsf{int}_\delta(y)$ contains at least two (integer) points (namely $\lceil x(1-\delta)\rceil$ and $\lceil x(1-\delta)\rceil + 1$). Invoking the Undominated Intersection Lemma (Lemma B.1), for any arbitrarily small ε, we can pick $\sigma_i \in \Delta(\mathsf{UDed}_i(\mathsf{int}_\delta(x)))$ and $\sigma'_i \in \Delta(\mathsf{UDed}_i(\mathsf{int}_\delta(y)))$ for every player $i \in N$ so that σ_i and σ'_i are ε-close in terms of both winning probability and price for any s_{-i}.

Now consider the outcome distribution $F(\sigma'_1, \ldots, \sigma'_n)$. (Note that, while the mechanism is deterministic, the randomness comes from the mixed strategy profile.) Since the good will be assigned with a total probability mass of 1, by averaging, there exists some player j such that $F^A_j(\sigma'_1,\dots,\sigma'_n) \le \frac{1}{n}$, i.e., player j receives the good with no more than $\frac{1}{n}$ probability. Without loss of generality, we assume that $j=1$.

According to the result of the Undominated Intersection Lemma, we know that $F^A_1(\sigma_1 \sqcup \sigma'_{-1}) \le \frac{1}{n}+\varepsilon$. If ε is small enough such that $\frac{1}{n}+\varepsilon<1$, then by the definition of expectation we know that there exists some *pure* strategy profile $s_1 \sqcup s'_{-1}$ that is in the support of $\sigma_1 \sqcup \sigma'_{-1}$ such that $F^A_1(s_1 \sqcup s'_{-1}) = 0$.

Now consider a "world" with δ-approximate-valuation profile K and true-valuation profile tv as follows:

$$\begin{aligned} K &\stackrel{\text{def}}{=} \big(\mathsf{int}_\delta(x), \mathsf{int}_\delta(y), \dots, \mathsf{int}_\delta(y)\big) \text{ and} \\ \mathsf{tv} &\stackrel{\text{def}}{=} \big(\lfloor(1+\delta)x\rfloor, \lceil(1-\delta)y\rceil, \dots, \lceil(1-\delta)y\rceil\big) \ . \end{aligned}$$

We have just shown that there exist some $s_1 \sqcup s'_{-1} \in \mathsf{UDed}(K)$ satisfying $F^A_1(s_1 \sqcup s'_{-1}) = 0$. Therefore,

$$\begin{aligned} \mathrm{SW}\big(\mathsf{tv}, F(s_1 \sqcup s'_{-1})\big) &= \lceil(1-\delta)y\rceil \le (1-\delta)y + 1 \le \frac{x(1-\delta)^2}{1+\delta} + 3 \\ &\le \frac{(1-\delta)^2}{(1+\delta)^2}\lfloor(1+\delta)x\rfloor + 4 \le \left(\frac{(1-\delta)^2}{(1+\delta)^2} + \frac{4}{B}\right)\lfloor(1+\delta)x\rfloor \\ &= \left(\frac{(1-\delta)^2}{(1+\delta)^2} + \frac{4}{B}\right)\mathrm{MSW}(\mathsf{tv}) \ . \end{aligned}$$

and thus the claimed inequality hods. ■

E Proof for Theorem 3

E.1 Proof of Statement 1

Looking for an optimal mechanism, we choose to restrict ourselves to consider only f-mechanism (recall Definition B.2), so as to leverage our Distinguishable Monotonicity Lemma (Lemma B.4) and simplify considerably our search. Indeed, this choice constrains the available strategies to coincide with all possible valuations, that is $\{0,\dots,B\}$, and our search is reduced to the the problem of finding a 1-DM winning-probability function f that ensures the desired fraction of the maximum social welfare at every possible tv_i in the integer interval $\{\min K_i,\dots,\max K_i\}$ (which, by the Distinguishable Monotonicity Lemma must contain *all* the undominated strategy profiles).

More concretely, we articulate the above high-level proof strategy in three conceptual steps.

Step 1

We start by putting forward a set of constraints for the winning-probability function f, and prove the they are sufficient for ensuring the target fraction of the maximum social welfare, *assuming* (for now) that the set of available strategies coincides with the set of all possible valuation profiles.

For every $\delta \in (0,1)$, define $D_\delta \stackrel{\text{def}}{=} \left(\frac{1+\delta}{1-\delta}\right)^2 - 1 = \frac{4\delta}{(1-\delta)^2} > 0$. We say that a winning-probability function f is *δ-good* if it is 1-distinguishably monotonic and:

$$\forall\, i \in N,\ \forall\, v \in \{0,1,\dots,B\}^N, \quad \sum_{j=1}^{n} f_j(v)v_j + D_\delta \cdot f_i(v)v_i \ge \frac{1}{n}\cdot v_i(n+D_\delta) \ . \tag{E.1}$$

We prove the following lemma:

Lemma E.1. *If f is δ-good, then the mechanism $M_f = (\Sigma, F_f)$ guaranteed by invoking the Distinguishable Monotonicity Lemma (Lemma B.4) with f has the following social welfare guarantee in undominated strategies: for every δ-approximate-valuation profile $K = (K_1, \dots, K_n)$, every $\mathsf{tv} \in K$, and every $v \in \mathsf{UDed}(K)$:*

$$\mathbb{E}\Big[\mathrm{SW}\big(\mathsf{tv}, F_f(v)\big)\Big] \geq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}\right) \mathrm{MSW}(\mathsf{tv}) \ .$$

Proof. For every player $i \in N$, let $x_i \in \mathbb{R}$ be such that $K_i \subset \mathsf{int}_\delta(x_i)$, and let $\mathsf{int}_\delta(x) = \mathsf{int}_\delta(x_1) \times \cdots \times \mathsf{int}_\delta(x_n)$.

Notice that $\mathrm{MSW}(\mathsf{tv}) = \mathsf{tv}_i$ for some i, so it suffices to prove the desired inequality by replacing the $\mathrm{MSW}(\mathsf{tv})$ by tv_i for each i. Then, re-ordering the four universal quantifiers of the lemma, we now proceed to show the following:

$$\forall i \in N, \quad \forall v \in \{0,1,\dots,B\}^N, \quad \forall \delta\text{-approximate } K \subset \mathsf{int}_\delta(x) \text{ s.t. } v \in \mathsf{UDed}(K), \quad \forall \mathsf{tv} \in K :$$

$$\sum_{j=1}^{n} \mathsf{tv}_j f_j(v) = \mathrm{SW}\big(\mathsf{tv}, F_f(v)\big) \geq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}\right) \mathsf{tv}_i \ . \tag{E.2}$$

Also, for every player $i \in N$: by Lemma B.4, we must have $(1-\delta)x_i \leq \min K_i \leq v_i \leq \max K_i \leq (1+\delta)x_i$, since f is 1-DM and $K_i \subset \mathsf{int}_\delta(x_i)$. Moreover, $\mathsf{tv} \in K$ indicates that $(1-\delta)x_i \leq \mathsf{tv}_i \leq (1+\delta)x_i$; combining these two we have

$$\frac{1-\delta}{1+\delta} v_i \leq \mathsf{tv}_i \leq \frac{1+\delta}{1-\delta} v_i \ . \tag{E.3}$$

In order to falsify Equation E.2, for a fixed i and v, we need to choose K and tv *adversarially* according to v. By analyzing the two sides of the inequality in Equation E.2, it is clear that choosing tv_j (for $j \neq i$) to be as small as possible is the "worst case", because tv_j only appears on the left and with a positive coefficient. However, depending on whether $f_i(v) - \big(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}\big)$ is non-negative or negative, the "worst case" choice of tv_i may be either as large as possible or as small as possible. For this purpose, we use the bounds of Equation E.3.

Based on the analysis above, by choosing the "worse case" choices of tv, the following inequalities become a "sufficient but not necessary" condition for implying the inequality of Equation E.2:

$$\forall\, i \in N,\ \forall\, v \in \{0,1,\dots,B\}^N, \quad \begin{cases} \sum_{j=1}^n \big(\frac{1-\delta}{1+\delta}\big) v_j f_j(v) \geq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}\right) \big(\frac{1-\delta}{1+\delta}\big) v_i \ , \text{ and} \\ \sum_{j=1}^n \big(\frac{1-\delta}{1+\delta}\big) v_j f_j(v) + \big(\frac{1+\delta}{1-\delta} - \frac{1-\delta}{1+\delta}\big) v_i f_i(v) \geq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}\right) \big(\frac{1+\delta}{1-\delta}\big) v_i \ , \end{cases}$$

or, equivalently,

$$\forall\, i\,, \forall\, v, \quad \begin{cases} \sum_{j=1}^n v_j f_j(v) \geq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}\right) v_i = \frac{n + D_\delta}{n} \cdot \frac{1}{D_\delta + 1} \cdot v_i \ , & \text{(E.4)} \\ \sum_{j=1}^n v_j f_j(v) + D_\delta \cdot v_i f_i(v) \geq \big((1-\delta)^2 + \frac{4\delta}{n}\big) \cdot \frac{1}{(1-\delta)^2} \cdot v_i = \frac{n + D_\delta}{n} v_i \ . & \text{(E.5)} \end{cases}$$

Note that Equation E.5 is the inequality required by the statement of the lemma; the other inequality, Equation E.4, we show is implied by Equation E.5. Indeed:

$$\sum_{j=1}^{n} v_j f_j(v) = \frac{1}{1 + D_\delta} \left(\sum_{j=1}^{n} v_j f_j(v) + D_\delta v_i f_i(v) \right) \geq \frac{1}{1 + D_\delta} \frac{n + D_\delta}{n} v_i \ ,$$

as desired. In sum, we have proved that as long as Equation E.5 is satisfied for all $i \in N$ and for any $v \in \{0,1,\dots,B\}^B$, the social welfare guarantees of Equation E.2 hold. $\square$

Step 2

We now must find a suitable δ-good winning-probability function f, letting the Distinguishable Monotonicity Lemma do "the rest of the work". Helpful in searching for such a function is the observation that a mechanism will have the toughest job maximizing social welfare whenever several players report high bids that are all very close. Because the players' valuations are only approximate, the mechanism will not be able to "infer" from such bids who is the player with highest true valuation. This suggests that assigning the good to a player chosen at random may be reasonable in such a situation. On the other hand, if players bids are not so clustered, then the mechanism should give a much higher probability mass to the highest bids, as lower bids are less likely to come from players with high true valuations. To achieve optimality, however, one must be much more careful in allocating probability mass, and some complexity should be expected.

We propose a specific winning-probability function $f^{(\delta)}$ (which depends on the approximation accuracy δ, which is indeed known to the mechanism designer), and prove that $f^{(\delta)}$ is well-defined, satisfies the constraints of Step 1, and also is 1-distinguishable monotonic.

Definition E.2. *For every $\delta \in (0,1)$, and let $D_\delta \stackrel{\text{def}}{=} \left(\frac{1+\delta}{1-\delta}\right)^2 - 1 > 0$. We define the function $f^{(\delta)} \colon [0,B]^N \to [0,1]^N$ as follows:*

- *for every $z = (z_1, \ldots, z_n) \in [0,B]^N$ such that $z_1 \geq z_2 \geq \cdots \geq z_n$, let n^* be the least index in N such that*
$$\forall\, i > n^*, \quad z_i \leq \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta}\ , \tag{E.6}$$
call players $1, \ldots, n^*$ the winners *and players $n^*+1, \ldots, n$ the* losers*, and then set*
$$f_i^{(\delta)}(z) \stackrel{\text{def}}{=} \begin{cases} \frac{1}{n} \cdot \frac{n+D_\delta}{n^*+D_\delta} \cdot \frac{z_i(n^*+D_\delta) - \sum_{j=1}^{n^*} z_j}{z_i D_\delta}, & \text{if } i \leq n^*, \\ 0, & \text{if } i > n^*; \end{cases} \tag{E.7}$$

- *for other z, define $f^{(\delta)}$ by extending it symmetrically: specifically, letting π be any permutation over the players such that $\pi(z) = (z_{\pi(1)}, \ldots, z_{\pi(n)})$ is non-increasing, we define $f_i^{(\delta)}(z) \stackrel{\text{def}}{=} f_{\pi(i)}^{(\delta)}(\pi(z))$.*

The definition of $f^{(\delta)}$ seems quite complicated, but the underlying intuition is not that obscure. Essentially, the first step in coming up for an educated guess for $f^{(\delta)}$ is to use symmetry to derive a candidate satisfying the constraints from Equation E.1. (The most natural such candidate is simply Equation E.7 with $n^* = n$.) Such a guess *almost* works, in the sense that it possesses all the properties that we want (the sum of the coordinates is at most 1 and it is 1-distinguishably monotonic), with the exception that it sometimes takes on negative values — and that is a problem as probabilities really must be non-negative.

The next step is thus to "patch" the guessed function by forcing non-negativity while maintaining all the other properties, and this is exactly where the idea of winners and losers comes in. Roughly, only players with sufficiently low reported valuations are at risk of a "negative probability" (and they are most likely to have low true valuations), so that we remove them from the auction altogether, and to preserve the other properties we need to re-weight the function, thereby obtaining Equation E.7. Thus, at high level, we simply keep removing players until all of the players are given non-negative probability (by virtue of being in the auction or having been thrown out). This idea essentially comes from another paper of the third author [CLS+10].

Having defined (and given some intuition for) our choice of $f^{(\delta)}$, we now turn to the task of proving that it satisfies all the properties that we want. However, the definition of $f^{(\delta)}$ is somewhat hard to work with, and proving properties about $f^{(\delta)}$ will require some work. For example, even establishing the simple property of monotonicity requires lots of care, because as a player's bid varies from 0 to B, the number of winners n^* also varies, thereby changing the expression for $f^{(\delta)}$, and thus it is not clear that the player's probability of winning does not go down.

Lemma E.3. *For $f^{(\delta)}$ from Definition E.2 the following properties hold:*

1. *The number of winners n^* is well defined.*
2. *For every $z = (z_1, \ldots, z_n) \in [0, B]^N$ such that $z_1 \geq z_2 \geq \cdots \geq z_n$, in addition to Equation E.6, we also have that,*
$$\forall\, i \leq n^*, \quad z_i > \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta} \ . \tag{E.8}$$
(And an analogous property holds for other z as $f^{(\delta)}$ is extended symmetrically.)
3. *The number of winners n^* is unique.*
4. *The function $f^{(\delta)}$ is a valid winning-probability function.*
5. *The function $f^{(\delta)}$ is monotonic and intergrable.*
6. *The function $f^{(\delta)}$ is 1-distinguishably monotonic.*
7. *The function $f^{(\delta)}$ is δ-good (it satisfies Equation E.1).*

Proof. We prove the statements one at a time:

(1) The requirement of Equation E.6 is always satisfied when $n^* = n$, so the set of possible n^*'s is not empty, and thus the smallest element of that set always exists.

(2) As z is assumed to be non-increasing, it suffices to prove Equation E.8 for $i = n^*$. And, indeed, by the minimality of n^* we know that if we attempt to choose $n^* - 1$, there exists some $j \geq n^*$ such that
$$z_{n^*} \geq z_j > \frac{\sum_{j=1}^{n^*-1} z_j}{n^* - 1 + D_\delta} \ ,$$
which, after rearranging, is equivalent to $z_{n^*} > \frac{\sum_{j=1}^{n^*} z_j}{n^* + D_\delta}$. Clearly, an analogous statement holds for other z, as $f^{(\delta)}$ is extended symmetrically, by re-labeling the indices of z to make it non-increasing.

(3) Suppose by way of contradiction that there exist two $n^\perp$ and $n^\top$ with $1 \leq n^\perp < n^\top \leq n$ satisfying Equation E.6; in particular, we have already established that they also satisfy Equation E.8. Now define $S^\perp \stackrel{\text{def}}{=} \sum_{j=1}^{n^\perp} z_j$, $S^\top \stackrel{\text{def}}{=} \sum_{j=1}^{n^\top} z_j$, $S^\Delta \stackrel{\text{def}}{=} S^\top - S^\perp$, and $n^\Delta \stackrel{\text{def}}{=} n^\top - n^\perp$.

By symmetry we only consider $z = (z_1, \ldots, z_n) \in [0, B]^N$ such that $z_1 \geq z_2 \geq \cdots \geq z_n$. By Equation E.6 and Equation E.8, for $i \in \{n^\perp + 1, \ldots, n^\top\}$,
$$\frac{S^\perp}{n^\perp + D_\delta} \geq z_i > \frac{S^\top}{n^\top + D_\delta} = \frac{S^\perp + S^\Delta}{n^\perp + n^\Delta + D_\delta} \ .$$

Averaging over all z_i for $i \in \{n^\perp + 1, \ldots, n^\top\}$, we get
$$\frac{S^\perp}{n^\perp + D_\delta} \geq \frac{S^\Delta}{n^\Delta} > \frac{S^\perp + S^\Delta}{n^\perp + n^\Delta + D_\delta} \ .$$

The second inequality of this last equation yields a contradiction with the first inequality in the same equation:

$$\frac{S^{\Delta}}{n^{\Delta}} > \frac{S^{\perp}+S^{\Delta}}{n^{\perp}+n^{\Delta}+D_\delta} \Leftrightarrow \frac{S^{\Delta}}{n^{\Delta}} > \frac{S^{\perp}}{(n^{\perp}+D_\delta)} \ .$$

(4) Substituting Equation E.8 into the definition of $f^{(\delta)}$ (defined in Equation E.7) we immediately have $f_i^{(\delta)}(z) \geq 0$ for each player i. Summing the $f_i^{(\delta)}$ up over all the players i we get:

$$\begin{aligned}\sum_{i=1}^{n} f_i^{(\delta)}(z) &= \frac{1}{n}\cdot\frac{n+D_\delta}{n^*+D_\delta}\cdot\sum_{i=1}^{n^*}\frac{z_i(n^*+D_\delta)-\sum_{j=1}^{n^*}z_j}{z_iD_\delta}\\ &= \frac{1}{n}\cdot\frac{n+D_\delta}{(n^*+D_\delta)D_\delta}\cdot\left(n^*(n^*+D_\delta)-\sum_{i=1}^{n^*}\sum_{j=1}^{n^*}\frac{z_j}{z_i}\right)\\ &\leq \frac{1}{n}\cdot\frac{n+D_\delta}{(n^*+D_\delta)D_\delta}\cdot(n^*(n^*+D_\delta)-n^*n^*) = \frac{n+D_\delta}{n}\cdot\frac{n^*}{n^*+D_\delta} \leq 1 \ .\end{aligned}$$

In particular, $f_i^{(\delta)}(z) \leq 1$ for each player i, as we have already established that $f_i^{(\delta)}(z)$ is non-negative. And thus, after symmetrically extending the above argument to all other z, we deduce that $f^{(\delta)}$ is a valid winning-probability function.

(5) For notational simplicity assume that $i = n$ (so that $z_{-i} = z_{-n}$) and also assume $z_1 \geq z_2 \geq \cdots \geq z_{n-1}$. (As usual, the other cases follow by symmetry by appropriate re-labeling.)

 – So define n' to be number of winners *when only considering the first* $(n-1)$ *values*, i.e., when only considering $z_1, \ldots, z_{n-1}$. We claim that:

$$z_n \leq \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta} \longrightarrow f_n^{(\delta)}(z) = 0 \ \text{(i.e., } n \text{ is a loser)} \tag{E.9}$$

$$z_n > \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta} \longrightarrow f_n^{(\delta)}(z) > 0 \ \text{(i.e., } n \text{ is a winner)} \tag{E.10}$$

The implication of Equation E.9 is clear, because the number of winners when player n is present is still n'. We now argue that the implication of Equation E.10 also holds, which is less of obvious.

So assume by way of contradiction that player n is a loser and yet $z_n > \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$. Let the current number of winners be n^* (i.e., when player n *is* present), so we know that $z_n \leq \frac{\sum_{j=1}^{n^*} z_j}{n^*+D_\delta}$; in particular, we must have that $n' \neq n^*$ (otherwise we are done, as there is already a contradiction). However, for this choice of n^*, we have that:

$$\forall\, i \in \{n^*+1, \ldots, n-1\}, \quad z_i \leq \frac{\sum_{j=1}^{n^*} z_j}{n^*+D_\delta} \ ,$$

But, this means that both n' and n^* are the number of winners when player n is absent, contradicting the uniqueness from Item 3. Thus, Equation E.10 also holds.

– Therefore, we only need to prove that given two $z_n^\perp, z_n^\top \in [0,B]^N$ such that $z_n^\top > z_n^\perp > \frac{\sum_{j=1}^{n'} z_j}{n'+D_\delta}$ we have that $f_i^{(\delta)}(z_{-n} \sqcup z_n^\top) \geq f_i^{(\delta)}(z_{-n} \sqcup z_n^\perp)$. Assume that when bidding $z_n^\top$ there are a total of $n^\top + 1$ winners (the first $n^\top$ players and player n), and bidding $z_n^\perp$ there are a total of $n^\perp + 1$ winners (the first $n^\perp$ players and player n). Then we claim that

$$n^\perp \geq n^\top \ . \tag{E.11}$$

Assume by way of contradiction that $n^\perp < n^\top = n^\perp + n^\Delta$. As before, let $S^\perp = \sum_{j=1}^{n^\perp} z_j$ and $S^\top = \sum_{j=1}^{n^\top} z_j = S^\perp + S^\Delta$. For every player $n^\perp \leq i < n^\top$:

$$\frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} < z_i \leq \frac{S^\top + z_n^\top}{n^\top + 1 + D_\delta} = \frac{S^\perp + S^\Delta + z_n^\top}{n^\perp + n^\Delta + 1 + D_\delta}$$

Averaging over all $n^\perp \leq i < n^\top$ we get:

$$\frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} < \frac{S^\Delta}{n^\Delta} \leq \frac{S^\perp + S^\Delta + z_n^\top}{n^\perp + n^\Delta + 1 + D_\delta}$$

but this is already a contradiction:

$$\frac{S^\Delta}{n^\Delta} \leq \frac{S^\perp + S^\Delta + z_n^\top}{n^\perp + n^\Delta + 1 + D_\delta} \Leftrightarrow \frac{S^\Delta}{n^\Delta} \leq \frac{S^\perp + z_n^\top}{n^\perp + 1 + D_\delta}$$
$$\Rightarrow \frac{S^\Delta}{n^\Delta} \leq \frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} \ ,$$

and thus Equation E.11 holds.

– Now we have established that $n^\perp \geq n^\top$. If $n^\perp = n^\top$ then $f_i^{(\delta)}(z_{-n} \sqcup z_n^\top) \geq f_i^{(\delta)}(z_{-n} \sqcup z_n^\perp)$ is immediately implied because they are using the same formula, and recall that Equation E.7 is non-increasing with respect to z_i.

If $n^\perp > n^\top$, let $n^\perp = n^\top + n^\Delta$, $S^\top = \sum_{j=1}^{n^\top} z_j$ and $S^\perp = \sum_{j=1}^{n^\perp} z_j = S^\top + S^\Delta$ as before. Then we average over all z_i for $n^\top < i \leq n^\perp$ we get:

$$\frac{S^\Delta}{n^\Delta} > \frac{S^\perp + z_n^\perp}{n^\perp + 1 + D_\delta} = \frac{S^\top + S^\Delta + z_n^\perp}{n^\top + n^\Delta + 1 + D_\delta} \Leftrightarrow \frac{S^\Delta}{n^\Delta} > \frac{S^\top + z_n^\perp}{n^\top + 1 + D_\delta} \tag{E.12}$$

Using that, we do the calculating:

$$\begin{aligned}
&f_n^{(\delta)}(z_{-n} \sqcup z_n^\top) - f_n^{(\delta)}(z_{-n} \sqcup z_n^\perp) \\
&= C_1 \cdot \Big(\frac{z_n^\top(n^\top + 1 + D_\delta) - S^\top - z_n^\top}{(n^\top + 1 + D_\delta) z_n^\top} - \frac{z_n^\perp(n^\perp + 1 + D_\delta) - S^\perp - z_n^\perp}{(n^\perp + 1 + D_\delta) z_n^\perp}\Big) \\
&= C_1 \cdot \Big(\frac{S^\perp + z_n^\perp}{(n^\perp + 1 + D_\delta) z_n^\perp} - \frac{S^\top + z_n^\top}{(n^\top + 1 + D_\delta) z_n^\top}\Big) \\
&= C_2 \cdot \Big((S^\perp + z_n^\perp)(n^\top + 1 + D_\delta) z_n^\top - (S^\top + z_n^\top)(n^\perp + 1 + D_\delta) z_n^\perp\Big) \\
&= C_2 \cdot \Big((S^\top + S^\Delta + z_n^\perp)(n^\top + 1 + D_\delta) z_n^\top - (S^\top + z_n^\top)(n^\top + n^\Delta + 1 + D_\delta) z_n^\perp\Big) \\
&= C_2 \cdot \Big(S^\top(n^\top + 1 + D_\delta)(z_n^\top - z_n^\perp) + S^\Delta(n^\top + 1 + D_\delta) z_n^\top - n^\Delta(S^\top + z_n^\top) z_n^\perp\Big) \\
&\geq C_2 \cdot \Big(S^\top(n^\top + 1 + D_\delta)(z_n^\top - z_n^\perp) + S^\Delta(n^\top + 1 + D_\delta) z_n^\top - n^\Delta(S^\top + z_n^\perp) z_n^\top\Big) \geq 0
\end{aligned}$$

Here we have used the fact of $z_n^{\top} - z_n^{\perp} \geq 0$ and $S^{\Delta}(n^{\top} + 1 + D_\delta) - n^{\Delta}(S^{\top} + z_n^{\perp}) > 0$ (by Equation E.12). This finishes the proof stating that $f^{(\delta)}$ is monotonic.

The integrability of $f^{(\delta)}$ is obvious, because $f^{(\delta)}$ is piecewise continuous, and there are at most n pieces, as the number of winners decreases when z_n increases (recall Equation E.11).

(6) Fix a player $i \in N$ and two distinct valuations $v_i, v_i' \in \{0, 1, \ldots, B\}$, and assume that $v_i < v_i'$. We have already established the monotonicity and intergrability of $f^{(\delta)}$, so that, to prove that $f^{(\delta)}$ is 1-DM, we only need to find a specific v_{-i} to make the integral positive (recall Definition B.3).

So define $v_{-i} \stackrel{\text{def}}{=} (v_i, v_i, \ldots, v_i)$, then:

- $f(v_i \sqcup v_{-i}) = \frac{1}{n}$ since there are n winners, all bidding the same valuation.
- $f(z \sqcup v_{-i}) = \frac{1}{nD_\delta}(D_\delta + n - 1 - \frac{v_i}{z}(n-1)) > \frac{1}{n}$, when $v_i < z \leq (1 + D_\delta)v_i$.

Here the upper bound on z is to make sure that the number of winners is still n. Notice that $f(z \sqcup v_{-i})$ is a function that is *strictly* increasing when z increases in such range, and therefore

$$\int_{v_i}^{v_i'} \big(f_i(z \sqcup v_{-i}) - f_i(v_i \sqcup v_{-i})\big)\, dz \geq \int_{v_i}^{\min\{v_i', (1+D_\delta)v_i\}} \big(f_i(z \sqcup v_{-i}) - f_i(v_i \sqcup v_{-i})\big)\, dz > 0 \ ,$$

as desired.

(7) Because $f^{(\delta)}$ is 1-DM according to Item 6, in order to prove that $f^{(\delta)}$ is δ-good, we only need to show that Equation E.1 holds. As usual, W.L.O.G. we assume $z_1 \geq z_2 \geq \cdots \geq z_n$.

We first observe that:

$$\begin{aligned}\sum_{i=1}^{n} f_i^{(\delta)}(z) z_i = \sum_{i=1}^{n^*} f_i^{(\delta)}(z) z_i &= \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} \cdot \sum_{i=1}^{n^*} \frac{z_i(n^* + D_\delta) - \sum_{j=1}^{n^*} z_j}{D_\delta} \\ &= \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} \cdot \left(\sum_{i=1}^{n^*} z_i\right) \ .\end{aligned}$$

For each loser i (i.e., with $i > n^*$), we know that

$$\sum_{j=1}^{n} f_j^{(\delta)}(z) z_j + D_\delta \cdot f_i^{(\delta)}(z) z_i = \sum_{i=1}^{n} f_i^{(\delta)}(z) z_i = \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} \cdot \left(\sum_{i=1}^{n^*} z_i\right) \geq \frac{1}{n} \cdot z_i \cdot (n + D_\delta) \ ,$$

as desired, where the last inequality is due to our choice of n^* (recall Equation E.6).

For each winner i (i.e., with $i \leq n^*$), we know that

$$\begin{aligned}\sum_{j=1}^{n} f_j^{(\delta)}(z) z_j + D_\delta \cdot f_i^{(\delta)}(z) z_i &= \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} \cdot \left(\sum_{i=1}^{n^*} z_i\right) + D_\delta \cdot f_i^{(\delta)}(z) z_i \\ &= \frac{1}{n} \cdot \frac{n + D_\delta}{n^* + D_\delta} z_i (n^* + D_\delta) = \frac{1}{n} \cdot z_i (n + D_\delta) \ ,\end{aligned}$$

again as desired. Notice that this is a generic argument for arbitrary $z \in [0, B]^N$, and when restricting $z = v \in \Sigma = \{0, 1, \ldots, B\}^N$, everything still holds; this finishes the proof that $f^{(\delta)}$ is δ-good.

$\square$

Step 3

Finally, we put all the pieces together together. The desired result will follow almost immediately by invoking Distinguishable Monotonicity Lemma, because it ensures that the undominated strategies are a subset of the possible valuation profiles.

Claim E.4. *For every $\delta \in (0,1)$, and consider the $f^{(\delta)}$-mechanism $M_{f^{(\delta)}} = (\Sigma, F_{f^{(\delta)}})$. For every δ-approximate-valuation profile $K = (K_1, \dots, K_n)$, every $\mathsf{tv} \in K$, and every $v \in \mathsf{UDed}(K)$,*

$$\mathbb{E}\Big[\mathrm{SW}\big(\mathsf{tv}, F_{f^{(\delta)}}(v)\big)\Big] \geq \frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} \cdot \mathrm{MSW}(\mathsf{tv}) \ . \tag{E.13}$$

Proof. By Lemma E.3, the function $f^{(\delta)}$ from Definition E.2 is a (well-defined) winning-probability function that is also δ-good. Therefore, by invoking Lemma E.1 with $f^{(\delta)}$, we deduce that the $f^{(\delta)}$-mechanism $M_{f^{(\delta)}}$ yields the target guarantee on social welfare in undominated strategies. $\square$

Finally, we remark that our proposed optimal probabilistic mechanism can indeed be computed efficiently (just like the second-price mechanism):

Lemma E.5. *The outcome function $F_{f^{(\delta)}}$ of the $f^{(\delta)}$-mechanism $M_{f^{(\delta)}}$ is efficiently computable.*

Proof. It suffices to show that both $F^A_{f^{(\delta)}} = f^{(\delta)}$ (the winning-probability function of the mechanism) and $F^P_{f^{(\delta)}}$ (the expected price function of the mechanism) are efficiently computable over $\{0, 1, \dots, B\}^N$.

First, we note that the winning-probability function $f^{(\delta)}$ is indeed efficiently computable, because the number of winners is between 1 and n and can be determined in linear time.

Next, we argue why the expected price function is also efficiently computable, which may not be so obvious as it is defined an integral of $f^{(\delta)}$ (see Definition B.2). So, without loss of generality, we show how to compute the expected price for player n as a function of v_n, i.e.,

$$f^{(\delta)}_n(v_{-n} \sqcup v_n) \cdot v_n - \int_0^{v_n} f^{(\delta)}_n(v_{-n} \sqcup z)\, dz \ .$$

Indeed, note that $f^{(\delta)}_n$ is a function that is piece-wisely defined according to different v_n, since different different values of v_n may result in different numbers of winners (i.e., values of n^*). So assume that $v_1 \geq v_2 \geq \cdots \geq v_{n-1}$, and let n' be the number of winners when player n is absent.

When $v_n \leq \frac{\sum_{j=1}^{n'} v_j}{n'+D_\delta}$, the proof of the monotonicity of $f^{(\delta)}$ (cf. Item 5 in Lemma E.3) implies that $f^{(\delta)}_n = 0$, so that integral below this line is zero.

When $v_n > \frac{\sum_{j=1}^{n'} v_j}{n'+D_\delta}$, one can again see from the proof of the monotonicity of $f^{(\delta)}$ that the number of winner players n^* is non-increasing as a function of v_n. Therefore, $f^{(\delta)}_n$ contains at most n different pieces and, for each piece, with n^* fixed, $f^{(\delta)}_n(v_{-n} \sqcup v_n) = a + b/v_n$ is a function that is inversely dependent on v_n so can be integrated symbolically. Therefore, the only question is how to calculate the pieces for $f^{(\delta)}_n$.

This is again not hard, by using a simple line sweep method. One can start from $v_n = \frac{\sum_{j=1}^{n'} v_j}{n'+D_\delta}$ and move v_n upwards. At any moment, one can calculate the earliest time that Equation E.8 is violated, and claim that another piece of $f^{(\delta)}_n$ is found. $\square$

E.2 Proof of Statement 2

Fix arbitrarily a number of players n, an approximation accuracy $\delta \in (0,1)$, and then a valuation bound $B \geq \frac{1}{\delta}$. We need to prove that:

For every (possibly probabilistic) *mechanism* $M = (\Sigma, F)$ *over* $\mathscr{C}^{\delta}_{n,B}$*, there exists a* δ*-approximate-valuation profile* K*, a strategy profile* $s \in \mathsf{UDed}(K)$ *and a true-valuation profile* $\mathsf{tv} \in K$ *such that*

$$\mathrm{SW}\big(\mathsf{tv}, F(s)\big) \leq \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} + \frac{4}{B} \right) \mathrm{MSW}(\mathsf{tv}) \ .$$

Define $x \stackrel{\text{def}}{=} B$ and $y \stackrel{\text{def}}{=} \lfloor \frac{x(1-\delta)+2}{1+\delta} \rfloor$. By our choice of B, we know that $x \geq y$. Therefore both $x, y \in \{0, 1, \ldots, B\}$. Furthermore, one can verify that $\mathsf{int}_\delta(x) \cap \mathsf{int}_\delta(y)$ contains at least two (integer) points (namely $\lceil x(1-\delta) \rceil$ and $\lceil x(1-\delta) \rceil + 1$). Invoking the Undominated Intersection Lemma (Lemma B.1), for any arbitrarily small ε, we can pick $\sigma_i \in \Delta(\mathsf{UDed}_i(\mathsf{int}_\delta(x)))$ and $\sigma'_i \in \Delta(\mathsf{UDed}_i(\mathsf{int}_\delta(y)))$ for every player $i \in N$ so that σ_i and σ'_i are ε-close in terms of both winning probability and price for any s_{-i}.

Now consider the outcome distribution $F(\sigma'_1, \ldots, \sigma'_n)$. Since the good will be assigned with a total probability mass of 1, by averaging, there exists some player j such that $F^A_j(\sigma'_1, \ldots, \sigma'_n) \leq \frac{1}{n}$, i.e., player j receives the good with no more than $\frac{1}{n}$ probability. Without loss of generality, we assume that $j = 1$.

According to the result of Undominated Intersection Lemma, we know that $F^A_1(\sigma_1 \sqcup \sigma'_{-1}) \leq \frac{1}{n} + \varepsilon$. Then by the definition of expectation we know that there exists some *pure* strategy profile $s_1 \sqcup s'_{-1}$ that is in the support of $\sigma_1 \sqcup \sigma'_{-1}$ such that $F^A_1(s_1 \sqcup s'_{-1}) \leq \frac{1}{n} + \varepsilon$.

Now consider a "world" with δ-approximate-valuation profile K and true-valuation profile tv as follows:

$$K \stackrel{\text{def}}{=} \big(\mathsf{int}_\delta(x), \mathsf{int}_\delta(y), \ldots, \mathsf{int}_\delta(y)\big) \text{ and}$$
$$\mathsf{tv} \stackrel{\text{def}}{=} \big(\lfloor (1+\delta)x \rfloor, \lceil (1-\delta)y \rceil, \ldots, \lceil (1-\delta)y \rceil\big) \ .$$

We have just shown that there exist some $s_1 \sqcup s'_{-1} \in \mathsf{UDed}(K)$ satisfying $F^A_1(s_1 \sqcup s'_{-1}) \leq \frac{1}{n} + \varepsilon$. Therefore, when choosing ε to be small enough, since we always have $\lceil (1-\delta)y \rceil < (1-\delta)y + 1$, we can make the following inequality go through:

$$\begin{aligned} \mathrm{SW}\big(\mathsf{tv}, F(s_1 \sqcup s'_{-1})\big) &\leq (\frac{n-1}{n} - \varepsilon) \cdot \lceil (1-\delta)y \rceil + (\frac{1}{n} + \varepsilon) \cdot \lfloor (1+\delta)x \rfloor \\ &< \frac{n-1}{n} \cdot (1-\delta)y + \frac{1}{n} \cdot \lfloor (1+\delta)x \rfloor + 1 \end{aligned}$$

We proceed this calculation and show that the social welfare satisfies the claimed inequality:

$$\begin{aligned} \mathrm{SW}\big(\mathsf{tv}, F(s_1 \sqcup s'_{-1})\big) &< \frac{n-1}{n} \cdot (1-\delta)y + \frac{1}{n} \cdot \lfloor (1+\delta)x \rfloor + 1 \leq \frac{n-1}{n} \cdot \frac{(1-\delta)^2 x}{1+\delta} + \frac{1}{n} \cdot \lfloor (1+\delta)x \rfloor + 3 \\ &< \frac{n-1}{n} \cdot \frac{(1-\delta)^2}{(1+\delta)^2} \lfloor (1+\delta)x \rfloor + \frac{1}{n} \cdot \lfloor (1+\delta)x \rfloor + 4 \\ &< \left(\frac{n-1}{n} \cdot \frac{(1-\delta)^2}{(1+\delta)^2} + \frac{1}{n} + \frac{4}{B} \right) \lfloor (1+\delta)x \rfloor \\ &= \left(\frac{n-1}{n} \cdot \frac{(1-\delta)^2}{(1+\delta)^2} + \frac{1}{n} + \frac{4}{B} \right) \mathrm{MSW}(\mathsf{tv}) \\ &= \left(\frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2} + \frac{4}{B} \right) \mathrm{MSW}(\mathsf{tv}) \ . \quad \blacksquare \end{aligned}$$

F Performance Diagrams

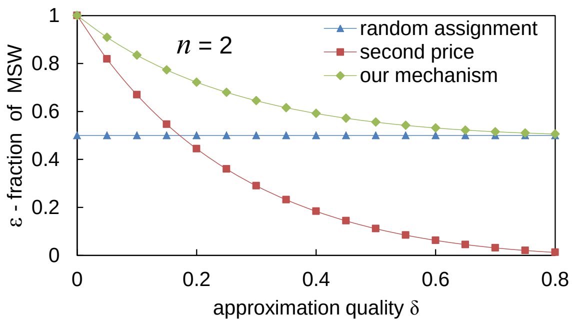

(a) With $n = 2$ players, the second-price mechanism performs *worse* than randomly assigning the good for $\delta > 0.18$.

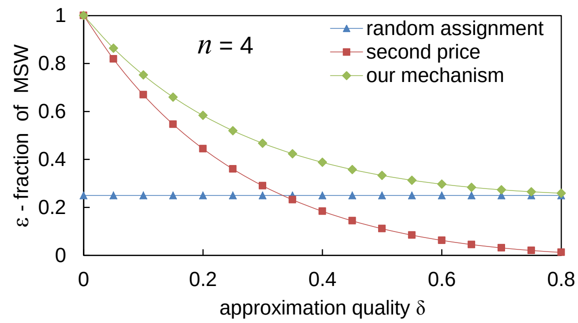

(b) With $n = 4$ players, the second-price mechanism performs *worse* than randomly assigning the good for $\delta > 0.34$.

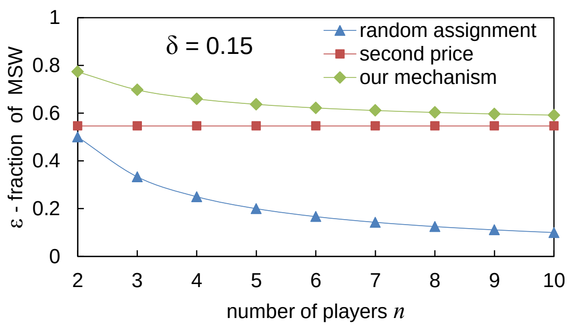

(c) With $\delta = 0.15$, the second-price mechanism always performs better than randomly assigning the good.

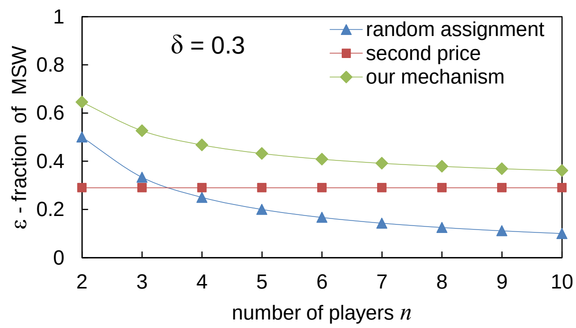

(d) With $\delta = 0.3$, the second-price mechanism performs *worse* than randomly assigning the good for $n = 2, 3$.

Figure 1: We compare the social welfare guarantees of randomly assigning the good ($\varepsilon = \frac{1}{n}$), the second-price mechanism ($\varepsilon = \frac{(1-\delta)^2}{(1+\delta)^2}$, see Theorem 2), and our optimal mechanism ($\varepsilon = \frac{(1-\delta)^2 + \frac{4\delta}{n}}{(1+\delta)^2}$, see Theorem 3). In (1a) and (1b) we compare ε versus δ, and in (1c) and (1d) we compare ε versus n. The green data, our mechanism, is always better (at times significantly) than the other two mechanisms.